

DIRECT PHOTON PRODUCTION IN RELATIVISTIC HEAVY-ION COLLISIONS

Paul Stankus

*Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831;
email: stankus@mail.phy.ornl.gov*

Key Words thermal, QCD, QGP, radiation, plasma

■ **Abstract** We examine the uses of direct photons in diagnosing the highly excited state of nuclear matter created in high-energy nuclear collisions. The traditional focus has been on direct photons as thermal radiation from the excited state, but we also explore the many other roles direct photons can play. We review experimental and theoretical techniques as well as the history of direct photon measurements in heavy-ion collisions and their interpretation.

CONTENTS

1. INTRODUCTION	517
1.1. What Is Meant by “Direct”?	518
1.2. Structure of This Review	520
2. SOURCES AND MOTIVATIONS	520
2.1. Prompt pQCD Processes	520
2.2. Thermal Radiation	524
2.3. Pre-Equilibrium Radiation	530
2.4. Parton-Medium Interactions	531
2.5. Other Electro-Weak Processes	531
3. EXPERIMENTAL TECHNIQUES AND RESULTS	533
3.1. The Statistical Subtraction Technique	533
3.2. Fixed-Target Results	535
3.3. Theoretical Interpretation	541
3.4. Other Techniques	545
3.5. Collider Experiments: First Results and Future Prospects	549

1. INTRODUCTION

Since its beginnings in the mid-1980s, a prime goal of the experimental program of colliding nuclei at high energies has been to create—and then study—highly excited, strongly interacting (i.e., hadronic or nuclear) matter over an extended region of space and time. Interest had been raised since the mid-1970s that such a region,

if it could be created in the laboratory, might show interesting properties, and these in turn would elucidate the nature of strong interactions in a nonperturbative regime. With a description in terms of QCD, widely taken to be the fundamental theory of strong interactions, the highly excited state of nuclear/hadronic matter is now generally referred to as a quark-gluon plasma, or QGP.

Well in advance of experiments, it was proposed that electromagnetic radiation would be a useful diagnostic of the QGP state. In 1976, for example, Feinberg (1) suggested that copious photon production would be a distinct feature of highly excited hadronic matter, and speculated that the number of photons could even exceed the number of hadrons in the final state after a QGP had cooled and decayed. The basic argument was straightforward: because electromagnetic interactions are much weaker than strong interactions, an hadronic system should be essentially transparent to photons, and so a large and long-lived QGP can radiate photons into the final state from throughout its entire volume and over its entire lifetime. Models of QGP radiation have been greatly refined over the last several decades, but the essential point has held true, that final-state photons are interesting because they can carry information from the entire history of a QGP's evolution (2, 3).

1.1. What Is Meant by “Direct”?

Observing direct radiation from a QGP state has been a central part of the motivation to measure photon production in relativistic heavy-ion collisions. However, direct photons of other kinds can also play important and interesting roles in diagnosing a QGP state. We explore the full palette of opportunities in this review. A sensible place to start is by defining the term direct photon, along with the related, but not synonymous, terms prompt photon, isolated photon, and thermal photon.

Sensibly enough, the term direct photon is used to indicate photons that emerge directly from a particle collision; they are distinguished from decay photons, which emerge as the daughters of long-lived secondaries which decay electromagnetically, such as $\pi^0 \rightarrow \gamma\gamma$ or $\Sigma^0 \rightarrow \Lambda\gamma$. For photons of interest produced in high-energy hadron collisions, i.e., hard gamma rays with energies $E > \sim 100$ MeV, the large majority will typically be decay photons, and so sifting out the direct photons is always an experimental challenge (see Section 3 for details). In Figure 1 we illustrate the basic suite of processes which can create direct photons in hadron collisions.

Direct photons can be created in hard scatterings of incoming partons. Panel A of Figure 1 shows the two perturbative QCD (pQCD) diagrams that produce final-state photons at lowest order: the gluon-photon Compton process (upper) and quark-antiquark annihilation (lower). A quark can also radiate photons after its initial hard scattering as part of the jet fragmentation process; this is depicted in the upper diagram of Panel B of Figure 1. To lowest order these are all the relevant processes for direct photon production in elementary hadron-hadron collisions (e.g., $p + p$ or $\pi + p$) taking place in vacuum. Their rates can in principle be calculated in pQCD and so they are sometimes called pQCD photons (also prompt photons; see below).

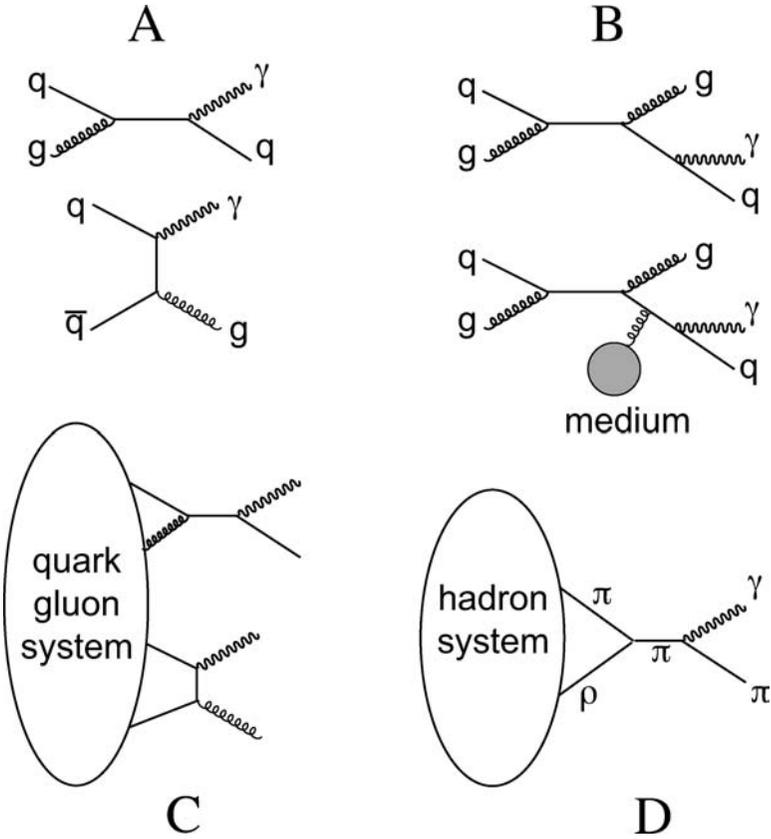


Figure 1 Illustration of the variety of processes which produce direct photons in hadron collisions. (A) Scattering between incoming partons. (B) Photons radiated by outgoing scattered partons, as part of the jet fragmentation process. (C) Scattering between quarks and gluons from a multi-collisional quark/gluon system. (D) Scattering between hadrons from a hadron system.

When a dense strongly interacting medium is present, however, more possibilities are opened up. If a scattered quark is traveling and fragmenting in a space-time co-occupied by such a medium, then photon radiation can be induced by parton-medium interactions, as in the lower diagram of Panel B of Figure 1. When a QCD medium is sufficiently hot, dense, and extended, then its constituents can scatter off each other and produce photons in the process. If the constituents are quarks and gluons, as shown in Panel C, then the lowest-level diagrams are the same as in Panel A. If the appropriate degrees of freedom in the medium are hadrons, then direct photons can be produced through interactions such as $\pi + \rho \rightarrow \pi + \gamma$, shown in Panel D. For a medium in thermal equilibrium, whether quark-gluonic, or hadronic (or other), its radiation is termed thermal photons. For a medium which is

self-interacting but has not yet reached equilibrium, its radiation could be termed pre-equilibrium photons, though the term is not yet in common use.

Now, a few more words on terminology. The final state of diagrams like those in Panel A of Figure 1 is a single direct photon recoiling against a jet. Photons in such a final state are termed isolated photons, or sometimes single photons, because the scattering produces no hadronic energy in a direction near the photon's. By contrast, photons from the processes in Panel B would always have hadrons accompanying the photon and are termed fragmentation photons or bremsstrahlung photons (and, in some older literature, anomalous direct photons). In high-energy physics it has been traditional to measure, and to treat theoretically, isolated photons as objects of primary interest, to the extent that isolated and direct are used almost interchangeably in some high-energy literature. In relativistic heavy-ion collisions, however, the final-state phase-space density of all hadrons is so high that measuring isolated photons is a practical impossibility, and so they are rarely discussed.

The term prompt photon is commonly heard, but its usage is not perfectly uniform. In the simplest definition, prompt refers to photons created on short time scales and so is synonymous with direct. In the more common usage, which we adopt in this review, prompt is used to distinguish photons from the processes in Panels A and B, which can be traced directly to incoming partons, from the processes in Panels C and D, which happen later in the evolution of a collision. This makes prompt photons equivalent to what we termed pQCD photons above. We can then state neatly that direct photons are the sum of prompt, pre-equilibrium, and thermal components.

1.2. Structure of This Review

All the varieties of direct photons described here have relevance to the diagnosis of a highly excited hadronic medium, or QGP, being created in relativistic heavy-ion collisions. In Section 2 we elaborate on each of these sources and its utility. In Section 3 we detail experimental techniques and review results of direct photon measurements in heavy-ion collisions, including both fixed-target and collider experiments. We also consider what lies ahead for the current programs.

Direct photon production in heavy-ion collisions is a broad subject with a substantial history, and we cannot hope to cover everything of relevance in a single article. The intent of this review is to be instructive, rather than encyclopedic, and we apologize to those whose work we could not find space to include or have overlooked in error. For more technically oriented reviews, see References 4–6. For useful, recent short summaries see References 7–9.

2. SOURCES AND MOTIVATIONS

2.1. Prompt pQCD Processes

To understand the roles that prompt pQCD photons can play in relativistic heavy-ion collisions, it is useful first to picture a collision's geometry. In Figure 2, we

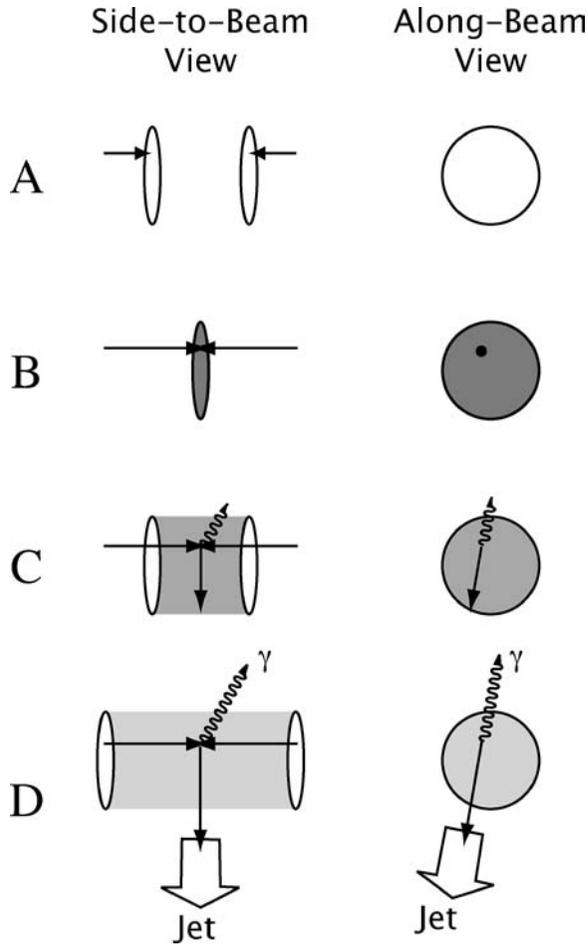


Figure 2 The geometry of a hard parton scattering to a jet+photon final state, taking place within a heavy-ion collision. (A) The incoming nuclei, before the collision, shown with Lorentz contraction along the beam direction. The arrows depict the paths of two partons within the nuclei. (B) At the point of collision the incoming nuclei overlap, and the partons undergo a hard scattering. The dark shade of the overlap region indicates a high energy density in an excited region. (C) As the remnants of the nuclei recede, the excited region expands and its energy density decreases, as indicated by the lighter shade. During the same interval the photon and scattered parton travel through the medium. (D) The medium continues to expand and cool. By this time the photon and the parton have exited the medium; in the final state the parton will be resolved as a jet of hadrons.

show a schematic picture of the time evolution of a heavy-ion collision in which the nuclei collide head-on and a hard parton-parton scattering takes place, which generates a jet and a photon. The scattering shown here is of the type in Panel A of Figure 1, where the jet (either quark or gluon) is balanced against a single photon. We can see that both the photon and the scattered parton, which later becomes a jet, must propagate through the QCD medium created by the collision.

In this example the photon is likely to emerge unscathed by the medium, while the parton/jet could have a significant interaction if the medium is sufficiently dense. Indeed, there is strong evidence that such dense media are created in heavy ion collisions at RHIC¹ energies. One of the most intriguing early results reported by the RHIC experiments was that of jet quenching: the yield of high- p_T hadrons in Au+Au collisions was greatly depleted, by a factor of up to 5, compared to a nominal scaling expectation (see below). Because jet fragmentation is thought to be the primary source of high- p_T hadrons, this depletion has been interpreted as strong parton/jet-medium interactions.²

The phenomenon of jet quenching has great potential as a diagnostic of the dense hadronic medium created in these collisions. Such interpretation is not without ambiguities, but several of these can be resolved with the help of direct photons.

2.1.1. PARTON FLUX, NUCLEAR STRUCTURE FUNCTIONS If we know the yield for some parton-parton hard scattering process in a p+p (or a nucleon+nucleon, hereafter N+N) collision, then how do we predict the yield for that process in a nuclear collision? The simplest approach, sometimes called point-like scaling or binary-collision scaling, follows from two assumptions: (a) Each nucleus is treated as a cloud of independent nucleons, each of which collides with nucleons in the opposing nucleus, possibly more than once. (b) Each of those N+N collisions is taken to be equivalent to an elementary N+N collision, even if one or both of the nucleons collides multiple times.³ Point-like scaling is the baseline assumption used in discussing yields of hard pQCD processes in nuclear collisions (“the nominal expectation” mentioned above).

¹The accelerator venues for heavy-ion collisions discussed in this review are RHIC (Relativistic Heavy-Ion Collider) and AGS (Alternating Gradient Synchrotron) at BNL, and SPS (Super Proton Synchrotron) and LHC (Large Hadron Collider) at CERN. See Section 3.2 for parameters and details.

²Although the effect at RHIC is dramatic, opinions differ as to whether jet quenching has been observed in hadron spectra from Pb+Pb collisions at CERN-SPS fixed-target energies, $\sqrt{s_{NN}} = 17$ GeV; see Reference 10.

³The justification for the latter assumption is twofold. First, hard parton scatterings are high- Q^2 processes that materialize quickly, before the soft processes that make up most of the N+N cross section, and so are unaffected by how many times a nucleon interacts. Second, because hard scatterings are very rare within each N+N interaction (hence point-like, meaning small cross section) it is reasonable to add their probabilities from different N+N collisions.

The dramatic departure from point-like scaling for high- p_T hadrons is the signature of jet quenching, and the most exciting interpretation for the violation is as an effect of a dense created medium on a scattered parton/jet. But point-like scaling could, in principle, be violated for other reasons as well. One frequently discussed possibility (11), for example, is that the combination of high $\sqrt{s_{NN}}$ and large nuclear size in RHIC Au+Au collision would cause the initial state of the incoming nuclei to be manifested as a color-glass condensate (CGC), a largely coherent state completely at odds with the assumption of independent nucleons described above. One effect of the CGC would be a substantial decrease in the flux of partons (compared with point-like scaling), which could lead to a decrease in the yield of high- p_T hadrons.

High-energy direct photons, which are produced primarily through processes like those shown in Figure 1A, allow us to separate the influence of initial-state effects (effects on partons before a hard scattering) from those of final-state effects (after the hard scattering, such as a created medium), because they are sensitive to the former but not the latter. As suggested in Figure 2, they can play the role of a control sample. A measurement such as the ratio between the yield of hadrons and the yield of pQCD photons, therefore, can indicate the presence of jet quenching directly without the need to invoke point-like scaling as a baseline (see Section 3.5.1).

At the same time, an absolute measurement of the yield of pQCD photons in relativistic heavy-ion collisions can serve to test the point-like scaling assumption for hard processes, effectively measuring the incoming parton flux in the incoming nuclei. A more exact statement is that they would allow the structure functions (equivalently, parton distribution functions) for nuclei to be compared with those for nucleons. This is one way that the presence of a CGC could be directly revealed, for example.⁴

These kinds of measurements are in principle very attractive for investigating the effects of nuclei on partons, but the interpretation of single inclusive direct photons in pQCD is not at present a settled matter theoretically (12), even for elementary collisions. It may be some time, then, before a reliable connection can be drawn between direct photons and parton distributions.

2.1.2. GAMMA-TAGGED JETS The fragmentation process can be described through jet fragmentation functions, which relate the energy of a scattered parton to the yields and spectra of the hadrons (and photons) which constitute its jet. Jet quenching, then, can be described (13) as a (severe) modification to these fragmentation functions brought on by the dense created medium.

In elementary hadron-hadron collisions, jets are identified as concentrations of hadronic energy in a narrow angular cone. The total energy of a jet can be

⁴Note that measurement of the Drell-Yan process $q\bar{q} \rightarrow l^+l^-$, which occurs through the production of a virtual photon, has the same advantages and can be seen as complementary to the measurement of single direct photons.

determined through calorimetry, and, in principle, fragmentation functions could be studied by comparing the spectra of particles in the cone to the total jet energy.⁵

This approach would be problematic in nuclear collisions, however, for two reasons. First, the calorimetric measurement of jet energies, and even the identification of jets as energy in a cone, will be greatly affected by the high general level of hadronic energy created in A+A collisions (and by the fluctuations of same). Second, without a good model for the jet quenching process, we cannot be sure that all of a parton's energy will even survive in hadrons correlated with the parton's direction. It is entirely possible that a significant portion of the quenched energy may be transferred to the dense medium and then redistributed over a wide range of angles.

If, however, a hard-scattered parton is balanced against a direct photon, as depicted in Figure 2, then the parton's energy can be reconstructed by measuring the photon. If we assume that the incoming partons have zero total momentum transverse to the beam axis, then the transverse momenta of the outgoing parton and photon must be equal and opposite. If the p_T of the photon is measured and the polar angle of the parton is known⁶ then the parton's total energy can be uniquely fixed, having been tagged by the photon.

This program for measuring medium effects on partons/jets opposite to a direct photon has been discussed widely [see for instance Wang and Huang (13), and also (14, 15)], and although it is in principle extremely promising, it does suffer from both experimental and theoretical difficulties. Experimentally, jet+photon final states are very rare among hard scatterings—the vast majority produce jet+jet final states—so this measurement will require much higher statistics than did the original quenching measurements. Also, the measurement requires the identification of direct photons on an individual, photon-by-photon basis, which can be challenging in the environment of a nuclear collision (see Section 3.4.2). Theoretically, there is strong evidence (16) that the hard scatterings which produce direct photons have nonzero total transverse momentum (intrinsic k_T) among their high- p_T outgoing particles, which complicates the energy tagging. However, these difficulties can all be reasonably controlled, and jet–direct photon correlations are actively being pursued, and their results eagerly awaited, at the time of this writing.

2.2. Thermal Radiation

2.2.1. THERMODYNAMICS OF QCD Early speculations that highly excited, strongly interacting matter might show interesting properties was put on a more quantitative footing with the advent of the lattice QCD computational technique (17). Roughly

⁵In practice fragmentation functions are usually studied in $e^+ + e^- \rightarrow \text{jet(s)}$ events, or deeply inelastic $l+p$ collisions, where the energies of scattered partons are known without measuring their daughter hadrons.

⁶Even though the full jet cannot be reconstructed, the direction of any (surviving) high-energy hadron is a reasonable approximation to the direction of its jet's parent parton.

speaking, lattice QCD simulates numerically a thermal ensemble of QCD field configurations. Ensemble averages can then reveal the thermodynamic properties—pressure, energy density, etc.—which a thermal system governed by the QCD Lagrangian (including, presumably, any locally equilibrated matter created in a heavy-ion collision) would have at any particular temperature.

This simple description belies the fact that lattice QCD calculations are highly technical and involve a number of approximations, and their results should not be mistaken for a closed-form description of QCD reality. There is no doubt, though, that the practice and technology of lattice QCD have made great advances in the last two decades, and they remain very active areas today. See (18–20) for recent general discussions of these techniques.

On very general grounds we can picture two limits for thermalized QCD: (a) At low temperatures, for example, well below the mass of a pion, we expect a QCD medium to look like a gas of (color-singlet) hadrons. (b) At very high temperatures, for example, tens of GeV or above, QCD color interactions weaken owing to asymptotic freedom, and a QCD medium should look like a nearly ideal gas of quarks and gluons. One might imagine that the change-over from one regime to the other could be gradual as a function of temperature. What lattice QCD reveals, however, is a sharp transition between two such regions, strongly suggestive of a phase transition at a critical temperature of about $T_C \simeq 170$ MeV. Results of this kind detailing a QCD phase transition have been a major source of motivation for the high-energy heavy-ion experimental program since its beginnings.

Figure 3 shows a modern calculation of such a transition, tracking the scaled energy density ε/T^4 versus temperature. For a relativistic Bose or Fermi gas, this normalized energy density is proportional to the number of different particle species, including the multiplicity of spin and color states; such a count is also (loosely) called the number of degrees of freedom in the system. The steep increase at $T = T_C$, followed by the plateau, can be interpreted⁷ as the opening from hadron to quark and gluon degrees of freedom. Whether or not this interpretation is completely justified (see Section 2.2.2 below), it has become standard to refer to these two phases of thermalized QCD as the QGP phase (or sometimes quark phase) and the hadron-gas or HG phase (sometimes also hadron-resonance gas, HRG).

Calculation of the rate of thermal direct photon emission from an equilibrated QGP and an equilibrated HG, which we discuss in Section 2.2.2 and Section 2.2.3, are the basic ingredients required to predict the spectra we would expect to see from a heavy ion collision. A full prediction also requires a detailed model of the collision's evolution; we return to this subject in Section 3.3.

⁷This interpretation is somewhat informal; more precisely, lattice QCD results describe two transitions in this region: deconfinement, which allows open color degrees of freedom, not just color singlets; and restoration of chiral symmetry, in which the quarks have no QCD-generated mass and their masses revert to their bare current (Higgs-generated) masses.

2.2.2. THERMAL QUARK-GLUON PLASMA RADIATION Calculating the emission of direct photons from a thermalized QGP, or HG, is an exercise in thermal field theory, whose techniques date back to the 1950s (21). The case of thermal QCD in particular has seen increased work and attention since the 1980s, in coincidence with the experimental program of high-energy nuclear collisions and the expectation of studying the QGP. Here we review briefly the basic theoretical ingredients relevant for real photon production and describe some representative calculations.

It should be noted that all calculations for the production of real photons are also immediately relevant to the production of virtual photons. Virtual photons are observable as continuum dilepton pairs, and the theory and observation of dileptons produced in heavy-ion collisions is a rich and important topic. Real and virtual photons should be thought of as two facets of the single subject of thermal electromagnetic radiation; and any calculation that compares with data for one of these observables—thermal direct photons or thermal dileptons—is, in principle, not complete until it is extended to compare with the other. We cannot even touch on the subject of dileptons in this review, but interested readers can start with (22) and references therein.

2.2.2.1. *Kinetic theory* The most immediately intuitive approach to calculating the rate of photon production in a thermal system is through kinetic theory. We identify a specific process by which the thermalized particles can produce a photon, and then fold the amplitude for that process in with the particles' thermal phase-space distributions.

For a process of the type $1 + 2 \rightarrow 3 + \gamma$, whose lowest-order diagrams are shown in Figure 1C, we denote $p^\mu = (E, \vec{p})$ as the four-momentum of the photon and $p_i^\mu = (E_i, \vec{p}_i)$ as the four-momenta of particles $i = 1, 2, 3$. If we define R_γ as the rate of photons produced per volume per time, then the contribution to the rate from this process can then be written as:

$$E \frac{dR_\gamma}{d^3p} = \frac{1}{2(2\pi)^3} \int \frac{d^3p_1}{2(2\pi)^3 E_1} \frac{d^3p_2}{2(2\pi)^3 E_2} f_1(E_1) f_2(E_2) (2\pi)^4 \delta^{(4)}(p_1^\mu + p_2^\mu - p_3^\mu - q_1^\mu) |\mathcal{M}|^2 \frac{d^3p_3}{2(2\pi)^3 E_3} [1 \pm f_3(E_3)]. \tag{1}$$

Here \mathcal{M} is the amplitude for the process; the $f_i()$ are the appropriate (Bose-Einstein or Fermi-Dirac) thermal phase space distributions of the strongly interacting particles; and the \pm implies either a Bose enhancement or a Pauli blocking (respectively) for the production of final-state particle 3.

The full rate is then obtained by summing over the rates for all possible photon-producing processes, and to be fully correct the amplitude for each process must be calculated over all orders of the strong coupling. (Because we only consider single photon emission, these are only valid to order e^2 in electromagnetic coupling.) An

advantage of the kinetic theory formulation, though, is that it is possible to make an approximate calculation perturbatively, starting with the simplest processes calculated at finite order. For this reason the kinetic approach has been widely used to predict photon rates from a QGP.

2.2.2.2. Photon self-energy A more formal, nonperturbative formulation is to recognize that the production rate of real, on-shell photons that escape a thermal system can be related (23) to the imaginary part of the in-medium, on-shell photon self-energy Π . For temperature T and photon energy E we have (24)

$$E \frac{dR_\gamma}{d^3 p} = -\frac{2}{(2\pi)^3} \text{Im}\Pi_\mu^{R,\mu} \frac{1}{e^{E/T} - 1}. \quad 2.$$

This relation is valid to order e^2 in the electromagnetic coupling, because it only considers the production of one photon, but in principle should be correct to all orders in the strong coupling.

The relation can be motivated, at a simple level, as follows (23, 25). The imaginary part of a particle's self-energy, or equivalently the imaginary part of its propagator, is associated with the disappearance of the particle's amplitude with time; in a zero-temperature vacuum it is directly related to the particle's decay rate. On-shell photons do not decay in vacuum, but in a medium they can disappear through processes which are the reverse of those shown in Figure 1C. But the rate at which a thermal system absorbs a particular particle at a specific energy and momentum must be balanced by the rate at which the system produces that same particle, so the connection between R_γ and $\text{Im}\Pi$ in Equation 2 emerges naturally.

The photon self-energy can be represented as an expansion in diagrams, and so evaluated perturbatively. One example at the two-loop level is shown in Figure 4; if the photon is in a thermal medium, the effects of the medium enter through the use of thermal propagators for the internal lines. The diagram expansion allows us to see a direct connection between the self-energy and kinetic theory representations of the photon rate: in Figure 4, the diagrams on either side of the cut line have exactly the same topology as the low-order photon production diagrams shown in Figure 1A and Figure 1C. A similar correspondence will hold for the higher-order photon production processes, and so perturbative evaluations in the two frameworks can be seen as essentially the same calculation.

2.2.2.3. Prototypical calculations Evaluating the rate R_γ using only the lowest-order photon production diagrams gives a result (4, 27, 28) whose essential features are shown here:

$$E \frac{dR_\gamma}{d^3 p} \propto \alpha \alpha_s T^2 e^{-E/T} \log \frac{ET}{k_c^2}, \quad 3.$$

where $\alpha = e^2/4\pi$ and $\alpha_s = g^2/4\pi$ are the electromagnetic (EM) and strong QCD couplings, and the proportionality factor will involve the number of quarks

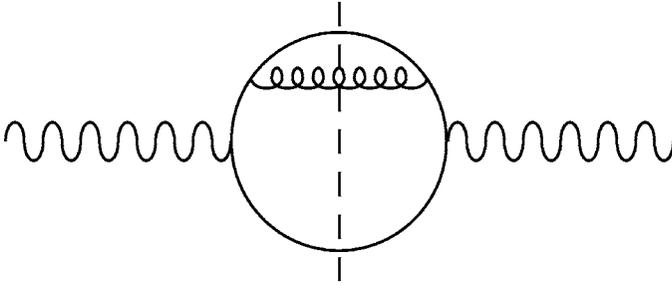


Figure 4 One QCD diagram that contributes to the photon self-energy. Each half of the diagram on either side of the cut (*vertical dashed line*) has the same topology as the lowest-order photon production diagrams in Figure 1A, C.

and the (squares of) their electrical charges. The result is infrared divergent in the mass of the intermediate quark line, and here is regulated with some infrared cutoff k_c .

The simplest choice for k_c would be the quark current mass. But this is unsatisfactory, in that the rate would be controlled and made finite by physics outside QCD, and no thermal medium information would enter into the quark propagator. Kajantie and Ruuskanen (26) had argued that the intermediate quarks should acquire an effective thermal mass in the medium; and the hard-thermal-loop (HTL) resummation of Braaten and Pisarski (29) showed that very soft modes of intermediate quarks and gluons would be suppressed in-medium, leading to a natural⁸ cutoff $k_c \sim gT$. The effective thermal mass was incorporated into the lowest-order calculation by Baier et al. (28), and the HTL cutoff was incorporated by Kapusta et al., who arrived at the result (27)

$$E \frac{dR_\gamma}{d^3p} = \frac{5}{9} \frac{\alpha\alpha_s}{2\pi^2} T^2 e^{-E/T} \log\left(\frac{2.912 E}{g^2 T}\right) \quad 4.$$

for two light quark flavors in the limit $E \gg T$; g and α_s are evaluated at the scale of the temperature.

The result of Equation 4 was used in many models of heavy-ion collisions to predict photon rates from the QGP phase. But it was subsequently pointed out by Aurenche et al. (30) that higher-order diagrams can contribute to the hard thermal photon rate at order $\alpha\alpha_s$ and so add significantly to the yield. Recently Arnold, Moore, and Yaffe have extended (31) the rate calculations to include the the Landau-Pomeranchuk-Migdal coherence effect and presented (32) complete leading-order results for QGP photon radiation.

⁸The HTL result is sometimes described loosely as in-medium quarks (and gluons) acquiring an effective mass. But these quarks do not acquire all the properties we would associate with mass—for example, they do not violate chiral symmetry—and so it is more accurate to describe the HTL result as a cutoff.

2.2.2.4. Photon rates in lattice QCD In Section 2.2.1 we described the very high temperature limit of QCD as approaching an ideal, weakly interacting gas of quarks and gluons. The graph in Figure 3 shows, however, that this limit is not reached even for $T \sim 4T_C$ —i.e., at temperatures approaching a GeV higher than are commonly expected to occur in heavy-ion collisions—since the energy density remains well below its ideal Stefan-Boltzmann value. This brings up the question of whether quarks and gluons are actually the appropriate degrees of freedom in a QGP phase near T_C for calculations in finite order. It has even been suggested (see 34 and references therein) that a variety of exotic massive states may appear in such a phase.

A complete calculation of thermal EM radiation with quarks and gluons to all orders would presumably account for the behavior of a QGP (or any other strongly interacting system) correctly. But because this cannot be accomplished in a perturbative expansion, an alternative approach is to calculate EM radiation within the framework of lattice QCD. The first lattice calculations of thermal dilepton rates have been carried out quite recently (33), and it is interesting to note that their results differ (35) from the perturbative calculation results at low photon energies. At the time of this writing, study of the subject is just beginning, and the validity of some of the techniques is being explored, but we are not aware of any fundamental reason why EM radiation should not eventually be on a par with other lattice QCD predictions.

2.2.3. THERMAL HADRON-GAS RADIATION In the early days of the relativistic heavy-ion collision program, there was an informal presumption that a QGP would produce a great deal more photon radiation than would a hadron gas at the same temperature. The (nonquantitative) reasoning ran along these lines: Because the quarks in a QGP are massless (due to the chiral symmetry restoration), they and their electrical charges will be present in large numbers, and because they are able to interact strongly (owing to deconfinement), we would expect to have a lot of charges being scattered and so a large amount of EM radiation—large, that is, relative nonquantitatively to a collection of massive hadrons bound up into color singlets. It came as something of a surprise, then, when Kapusta, Lichard, and Seibert (KLS) made the first serious calculation (27) of photon radiation from a thermal hadron gas and announced that at relevant temperatures (~ 200 MeV), the HG and QGP would produce very similar spectra of radiated photons.

The technology for calculating thermal radiation from a hot HG is the same thermal field theory as discussed above in Section 2.2.2. The same two formulations, kinetic theory and photon self-energy, will apply, but now with hadron processes and diagrams like those in Figure 1D. The programs are in principle very similar, but there are two factors that further complicate the thermal hadron gas case.

First, hadrons are more complicated objects than quarks or gluons. We know that the natural lineshape of the ρ resonance is broad, with a mass of 770 MeV and a width of 150 MeV. Further, the properties of the ρ are expected to change significantly in a dense hadronic medium in ways that are not straightforward to calculate. These complications—which are not limited just to the ρ —must be taken

into account in order to make any realistic estimate of electromagnetic radiation from a hadron gas.

Second, while there is only one gluon and two (or three) flavors of light quark, there are many different hadrons, and it is not always immediately apparent which might be important in calculating photon production. For example, the KLS calculation included processes involving π , ρ , and η mesons, and the decay $\omega \rightarrow \pi^0 \gamma$. But shortly thereafter Xiong, Shuryak, and Brown pointed out (36) that processes involving the $A_1(1260)$ resonance provided new channels, which would significantly raise the rate for photon production. Later refinements included: processes with strange mesons and resonances; processes with ω 's in the t -channel; and the effects of nonzero net baryon density. So the fidelity of hadron gas calculations depends on the level of sophistication one is able to bring to them, even in something as basic as the choice of which degrees of freedom to consider.

Recently, reasonably state-of-the-art calculations for radiation from a thermal hadron gas were presented (37, 38) by Turbide, Rapp, and Gale, considering all the effects listed above and including a massive Yang-Mills model of hadron interactions. Their conclusion, interestingly, is essentially the same as that of KLS, that the thermal photon spectra from a hadron gas are very similar to that from a QGP at the same temperature; Steffen and Thoma (39) reached a similar conclusion. This outcome is somewhat ironic, perhaps, after a decade's worth of increased sophistication in both HG and QGP calculations.

2.3. Pre-Equilibrium Radiation

As mentioned in Section 1, between the prompt source and the thermal source lies the pre-equilibrium source of photons, referring to processes that are not associated with one particular parton scattering but that take place before local thermal equilibrium obtains. Pre-equilibrium photons are one of the few observables that could provide direct information on this not-well-understood—but potentially novel and interesting—stage of a heavy-ion collision.

The simplest version of a pre-equilibrium state is one that is close to, but not quite in, local equilibrium; an example would be a system which is equilibrated kinetically but not chemically. In the “hot-gluon” scenario (40), for instance, gluons equilibrate quickly in a nuclear collision but the quarks lag in coming up to their equilibrium density; as a result (40, 41) EM radiation is less intense than the radiation from fully equilibrated quarks and gluons at the same (kinetic) temperature. However, at fixed energy density, an underpopulation of quarks is compensated by an increase in the kinetic temperature, making the photon rates similar to those from true equilibrium (42).

Descriptions of such near-equilibrium states offer no insight on what processes bring about kinetic equilibrium initially. The subject of pre-equilibrium photon radiation has been treated in a number of approaches, including parton cascades (43), microscopic transport (44), and nonequilibrium field theory (45, 46), and found to be a potentially significant contribution to the direct photon rate. However, the overall subject of the pre-equilibrium stage of a heavy-ion collision, and how it evolves to local equilibrium, is decidedly underexplored at present.

2.4. Parton-Medium Interactions

The phenomenon of jet quenching in heavy-ion collisions (see Section 2.1) provides strong evidence for, and in principle a diagnostic of, a dense medium created in those collisions. Direct photons associated with jets (fragmentation or bremsstrahlung photons) offer a new window into parton-medium interactions.

Photons appear, at some level, alongside hadrons when a scattered parton fragments into a jet, as in the vacuum diagram of Figure 1B. If partons lose energy in the medium before they fragment, then the population of high- p_T fragmentation photons will be suppressed along with high- p_T hadrons. Jeon et al. (47) estimated that this could lead to as much as a 20–40% reduction in the prompt photon spectrum from a RHIC Au+Au collision for $p_T > 3$ GeV/c. Zakharov argued (49), by contrast, that a dense medium acting on an energetic scattered quark would induce additional photon bremsstrahlung as in the in-medium diagram of Figure 1B. The net effect would be a substantial increase in prompt photons in RHIC and LHC collisions in a range $5 < p_T < 15$ GeV/c.

Fries, Müller & Srivastava (48) noted that scattered fast quarks could create photons by annihilating on antiquarks in the created medium, or through a Compton scattering with a gluon in the medium. These interactions would provide a direct measure of the density of quarks and gluons in the medium. The same mechanisms could also lead to the production of high- p_T , high-mass dilepton pairs (50). Fries et al. predicted that this process would be the dominant source of direct photons for $p_T < 6$ GeV/c in RHIC and LHC collisions.

These pictures of enhanced photon production from parton-medium interactions imply an interesting experimental signature: As the parton travels through a longer or denser medium, the number of prompt photons will increase while the number of high- p_T hadrons decreases. This anti-correlation could be a unique way to observe these medium-induced photons apart from other kinds of direct photons. Angular correlations among photon-hadron pairs are another natural way to observe jet fragmentation photons, and results along these lines at RHIC are eagerly awaited.

2.5. Other Electro-Weak Processes

As mentioned above, the production of real photons in a QCD system is intimately related to the production of virtual photons; and with only a little latitude, we can consider the entire electro-weak sector as part of the same subject. Taking that latitude, we here make some brief mention of how the production of heavy electro-weak W^\pm and Z bosons may be of interest in relativistic heavy-ion collisions at collider energies.

2.5.1. PARTON DISTRIBUTIONS AND JET TAGGING Beam energies at RHIC and LHC will be high enough to create, for the first time ever, real (i.e., s -channel) W s and Z s in collisions involving nuclei in the laboratory. For a $2 \rightarrow 1$ parton process, the mass of the created particle is related to the Bjorken x of the partons by $M^2 = x_1 x_2 s$. Heavy electroweak boson production at RHIC, then, is a

fairly high- x process with average $\sqrt{x_1 x_2} \sim 0.4$, and at LHC a more modest $\sqrt{x_1 x_2} \sim 0.015$.

Partonic cross sections for W and Z production are straightforward to calculate and so measuring their production in collisions with nuclei provides a measure (51) of nuclear effects—shadowing at LHC, Fermi motion at RHIC—on quark and antiquark distributions at very high $Q^2 \sim 10^4$ GeV² scales, much higher than have ever been explored in nuclei before. These nuclear effects can be studied in either $p+A$ or $A+A$ systems; but $A+A$ has the advantage that their geometric dependence, i.e., variation with nuclear impact parameter, can be more readily mapped out, which will provide important information as to their origin.

Heavy electroweak bosons can also play roles in the study of jet production and jet quenching in nuclear collisions. It has been suggested (52) that because Z s can be reconstructed with great accuracy through their dilepton channel, very high- p_T Z s could serve to tag high-energy jets in a manner similar to the photon tagging of jets discussed in Section 2.1.2. Also, W s and Z s will provide a steady, calculable⁹ supply of pairs of high-energy quark/antiquark jets through their $q\bar{q}$ decay. Departures from the expected appearance of these jet pairs will provide a measure of the dense medium's effect on (anti)quark jets separately from the more common gluon jets.

2.5.2. DENSE-MEDIUM Z MODIFICATION The heavy electro-weak bosons have the potential to serve as dedicated early-time probes in a heavy-ion collision. A Z will be formed quickly in an initial hard scattering ($\hbar/M_Z \sim 10^{-3}$ fm/ c) and live for only a short time ($\hbar/\Gamma_Z < 0.1$ fm/ c), and so will experience only the earliest equilibrated and pre-equilibrium stage of the collision. Given this, it is interesting to ask whether the properties of the Z should be measurably¹⁰ modified in a hot/dense QCD medium.

Changes in the mass and/or decay width of the Z follow from how its propagator, or self-energy, is modified in the thermal medium. This is essentially the same calculation described in Section 2.2.2 for photons, and it can be calculated perturbatively with diagrams like the diagram in Figure 4, with a Z line in place of the photon line. (Semi-classically we can think of the Z as undergoing collisional broadening by the medium; the connection between thermal self-energies and multiple scattering in a medium has been discussed (54).)

Kapusta & Wong (53) calculated these modifications at the two-loop level and concluded that the effects would be small. For temperature $T = 1$ GeV, the

⁹In LHC collisions we would expect pairs of jets with invariant mass equal to either M_W or M_Z , whereas at RHIC the antiquark distribution is falling so steeply at high x that the W and Z lineshapes are greatly distorted on the low-mass side. But due simply to the rates that can be expected this measurement is probably practical only at the LHC anyway.

¹⁰Although the medium effects on W^\pm and Z will be very similar, we focus here on the Z since its mass and decay width can be measured much more accurately through its dilepton decay channel.

expected increase in the Z width would be on the order of 1 MeV, and the change in mass even smaller, compared with the natural (vacuum) width of 2.5 GeV. Technical objections to this calculation were raised by Aurenche et al. over the handling of divergences (see Reference 55 for the comment and Reference 56 for the reply). The subject of divergences in thermal boson self-energies was examined in detail by Majumder and Gale (57), who presented results that differed from those in (53) by logarithmic factors. However, there is no practical reason at present to doubt Kapusta and Wong's basic conclusion that effects on the Z from a thermalized QGP that can be produced in heavy-ion collisions will be unobservable.

It remains an open question, however, whether the pre-equilibrium medium might have a larger, measurable effect on the electro-weak bosons—although this is primarily because there is currently no good theory of what physics bridges the initial scatterings and the initial locally thermalized state (even the use of the word “medium” for the pre-equilibrium stage may be an overstatement; see also Section 2.3). We are not aware of any calculations along these lines¹¹, but it remains true on general grounds that Z would constitute a focused probe of the very earliest stages of a heavy-ion collision.

3. EXPERIMENTAL TECHNIQUES AND RESULTS

3.1. The Statistical Subtraction Technique

The statistical subtraction technique embodies the very definition of direct photons: (a) measure the spectrum of all photons $\gamma^{Incl}(p_T)$ inclusively; (b) calculate the spectrum of photons $\gamma^{Decay}(p_T)$ produced in hadronic decays; and (c) subtract the decay contribution from the inclusive spectrum, and the difference must be the direct spectrum as

$$\gamma^{Dir}(p_T) = \gamma^{Incl}(p_T) - \gamma^{Decay}(p_T). \quad 5.$$

Although this program is in principle well defined, the experimental practice has challenges at each step. (Here we can only scratch the barest surface of the experimental details; see Reference 58 for an exhaustive discussion.)

3.1.1. PHOTON DETECTION There are two general approaches to detecting and reconstructing final-state photons. First, their energy and position can be measured in a transversely segmented electromagnetic calorimeter. Second, they can undergo conversion in a thin layer of material into an electron-positron pair, which are then tracked as charged particles. The first method works well for photons with high lab energy, because the fractional resolution of calorimeters improves with energy; but it will become difficult, typically, for photons with lab energies below 0.5–1.0 GeV.

¹¹This does not mean such calculations could not be attempted today. The possibility that pre-equilibrium gluons could affect the Z propagator could be addressed in a parton cascade framework, for example.

The conversion method will provide better resolution at these low photon energies, but will lose resolution with increasing lab energy. Additionally, detection through conversion will suffer a loss of efficiency because the converter material layer must be kept thin, on the order of a few percent of a radiation length, so as not to degrade the charged particle tracking. Overall the segmented-calorimeter approach is more common in photon detection experiments.

3.1.2. MESON RECONSTRUCTION In all hadron collisions the primary source of decay photons is from the decay of π^0 's, which have a 99% branching ratio to two photons and typically produce 85–90% of all decay photons. Other important sources include the decays of the η (5–10% of decay photons) and η' and ω (on the order of 1% each). For the event sample of interest, the spectra of all these mesons must be measured, or very well estimated, in order to calculate the spectrum of decay photons.

The most common method of reconstructing π^0 s is by detecting the photons of their primary decay $\pi^0 \rightarrow \gamma\gamma$. Photon pairs with a (combined) p_T in some specified range are binned in pair invariant mass $m_{\gamma\gamma}$, and π^0 s appear as a peak at $m_{\gamma\gamma} = m_{\pi^0}$ on top of a combinatorial background. The signal-to-background ratio (S/B) in the peak region is reduced at low p_T and in more central nuclear collisions with high photon multiplicities; yields of π^0 peaks have been reliably extracted for S/B below 1%. Systematic uncertainties on π^0 peak extraction are usually¹² the dominant source of error in direct photon measurements using the subtraction method.

To keep relative normalization uncertainties at an acceptable level, the spectrum of π^0 s should be measured in the same experiment and over the same event sample as the inclusive photons. The heavier mesons can in principle be reconstructed in a similar way, but their statistical significance will be reduced by their lower yields and low branching ratios to all-photon final states. A more precise estimate of their spectra can be made by assuming some scaling behavior¹³ relative to the π^0 spectrum and then normalizing to meson/ π^0 ratios measured at a few points in p_T or taken from elementary hadron collisions.

3.1.3. DECAY SUBTRACTION Once the spectra of the contributing mesons are known, the spectrum of their decay photons $\gamma^{Decay}(p_T)$ can be calculated straightforwardly. With the inclusive photon spectrum $\gamma^{Incl}(p_T)$, one can then in principle calculate $\gamma^{Dir}(p_T)$ as per Equation 5.

However, it is advantageous for a number of reasons to work instead with the ratio of spectra $\gamma(p_T)/\pi^0(p_T)$ (or γ/π^0 for short). A number of systematic

¹²There are approaches to measuring the spectra of π^0 s by inferring them from the spectra of charged pions; see Section 3.2 for examples. But these methods generally do not yield the same sensitivity in direct photon measurement as does the reconstruction of π^0 s.

¹³A standard assumption is that the meson spectra have the same shape when plotted against the transverse mass $m_T = \sqrt{p_T^2 + m^2}$.

uncertainties can be reduced or eliminated this way; for example, the effects of uncertainty in the overall, linear scale of the reconstructed energy will cancel directly in this ratio, assuming that the photons and π^0 s are measured in the same detectors (note that the effects from any potential nonlinearity in the detectors' energy response can still be important). Also, γ/π^0 is more convenient to examine, because although the individual spectra may vary by several orders of magnitude over a few GeV/c range in p_T , the γ/π^0 ratio will typically vary by substantially less than one order of magnitude. In fact, if the π^0 spectrum follows a power law, i.e., $\propto (p_T)^{-n}$ for some power n , then at high $p_T \gg m_\pi$ the γ/π^0 ratio will simply go to a constant.

The earlier generation of direct photon analyses (see Reference 60, for example) accordingly focused on the γ/π^0 ratio as the relevant observable and reconstructed the direct photon spectrum as

$$\gamma^{Dir}(p_T) = \left[\frac{\gamma^{Incl Meas}(p_T)}{\pi^0 Meas(p_T)} - \frac{\gamma^{Decay Sim}(p_T)}{\pi^0 Sim(p_T)} \right] \pi^0 Meas(p_T). \quad 6.$$

Here we have explicitly distinguished between several spectra: $\gamma^{Incl Meas}$ and $\pi^0 Meas$ are the measured inclusive spectra, with their statistical fluctuations. The measured $\pi^0 Meas$ is then represented by $\pi^0 Sim$, typically a smoothed functional form from which the spectra of the heavier mesons and the decay photons $\gamma^{Decay Sim}$ can then be calculated.

A later refinement (63) was to focus on the so-called double ratio observable, here termed R :

$$R(p_T) \equiv \frac{\gamma^{Incl Meas}(p_T)/\pi^0 Meas(p_T)}{\gamma^{Decay Sim}(p_T)/\pi^0 Sim(p_T)}. \quad 7.$$

The double ratio $R(p_T)$ will now, in principle, be a measure of $\gamma^{Incl}(p_T)/\gamma^{Decay}(p_T) = 1 + \gamma^{Direct}/\gamma^{Decay}$. The presence of direct photons will be revealed as a departure from $R = 1$, whereas the null result of no detectable direct photons will present itself as R being consistent with 1 over all p_T .

An example of the latter case is shown in Figure 5. Direct photons are not detectable from these data, meaning that one cannot set a nonzero lower limit on their yield at any p_T . However, it should be noted that it is always possible to set an upper limit on the yield at any p_T . (See Figure 7 for the limits corresponding to Figure 5, and the discussion in Section 3.2.) An example of double ratios which rise significantly above 1 are shown in Figure 6, and the corresponding direct photon spectra is shown in Figure 12.

3.2. Fixed-Target Results

The experimental program of high-energy heavy-ion collisions can be divided into three successive stages:

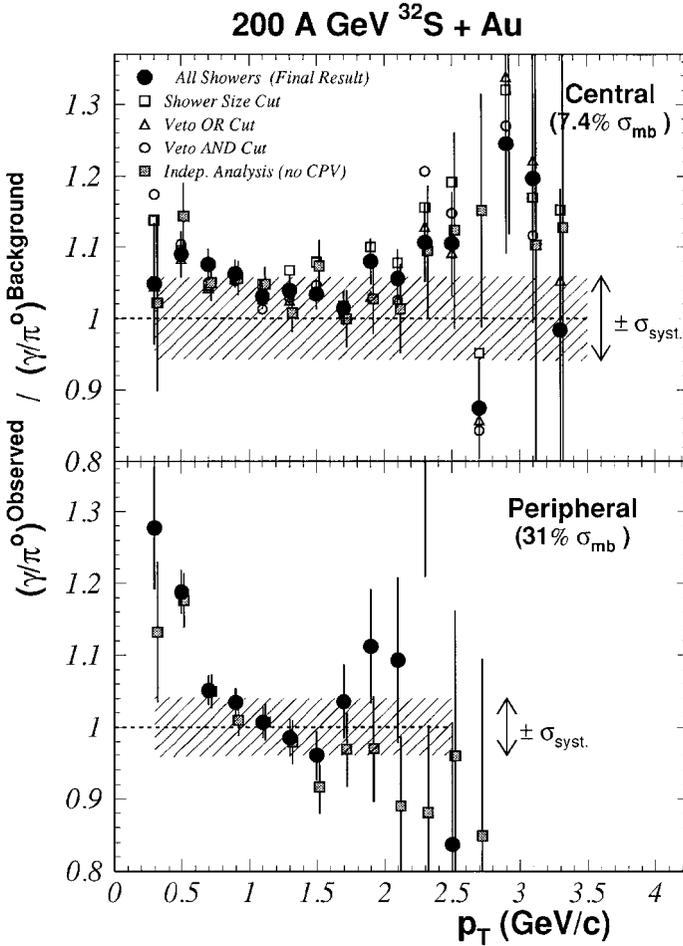


Figure 5 The double ratio (see text) versus p_T for central and peripheral S+Au collisions at $\sqrt{s_{NN}} = 19$ GeV at the CERN-SPS (63). In both event samples the double ratio is consistent, to within errors, with 1 at all p_T . This means that no lower limit can be set on the yield of direct photons from these data; the upper limits are shown in Figure 7.

1. Fixed-target experiments with beams of lighter nuclei,¹⁴ starting in the mid-1980s: ¹⁶O and ²⁸Si beams were accelerated at the BNL-AGS to 14.6 AGeV/c ($\sqrt{s_{NN}} = 5.4$ GeV), and ¹⁶O at 60 AGeV/c ($\sqrt{s_{NN}} = 11$ GeV) at the CERN-SPS followed by ¹⁶O and ³²S beams at 200 AGeV/c ($\sqrt{s_{NN}} = 19$ GeV).

¹⁴Linguistic purists may object to hear a light nucleus being referred to as a heavy ion; but the term “heavy ion” is an archaic one, used to denote any largely or fully ionized nucleus.

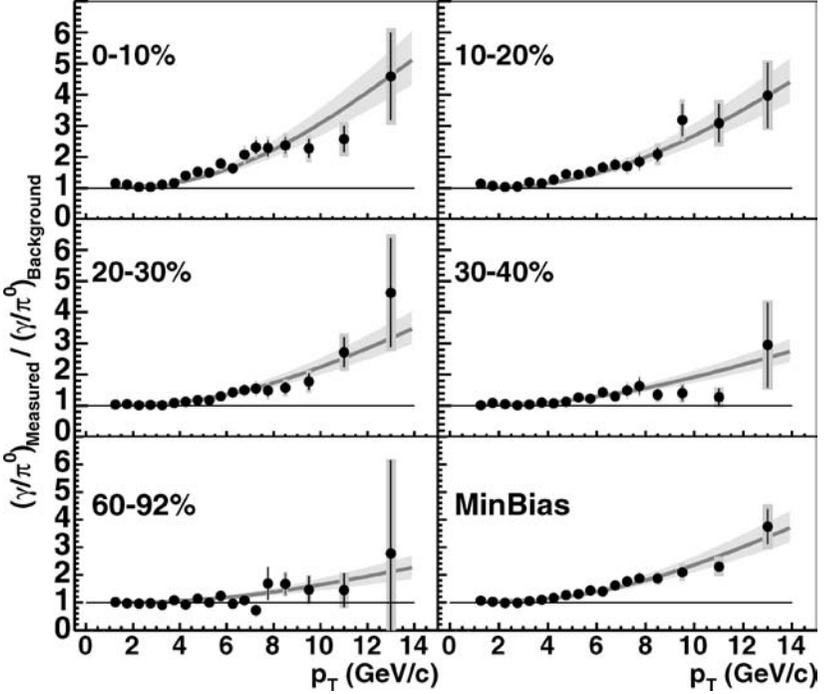


Figure 6 The double ratio versus p_T for various centralities in Au+Au collisions at RHIC, $\sqrt{s_{NN}} = 200$ GeV, as measured by the PHENIX experiment (93) and compared to NLO pQCD calculations. The resulting spectra are shown in Figure 12.

2. Fixed-target experiments with beams of heavy nuclei, starting in the early 1990s: ^{192}Au at 11.7 AGeV/c ($\sqrt{s_{NN}} = 4.8$ GeV) at the AGS, and ^{208}Pb at 160 AGeV/c ($\sqrt{s_{NN}} = 17$ GeV) at the SPS.
3. Collider experiments, which started in the year 2000 with Au+Au at BNL-RHIC at energies of $\sqrt{s_{NN}} = 130$ GeV, followed by Au+Au at $\sqrt{s_{NN}} = 200$ GeV and then a number of other combinations of nuclei and energies. Pb+Pb collisions at the CERN-LHC at $\sqrt{s_{NN}} = 5.5$ TeV should be available before the end of this decade.

In this section we summarize the results of direct photon measurements in SPS fixed-target experiments (none were reported in AGS experiments). Recent RHIC results are reviewed in Section 3.5.1, and future LHC prospects are discussed in Section 3.5.2.

3.2.1. LIGHT-NUCLEUS BEAMS All the direct photon results from SPS experiments described in this section used some form of the statistical subtraction method (see Section 3.1); other more recent SPS results obtained at very low p_T are discussed

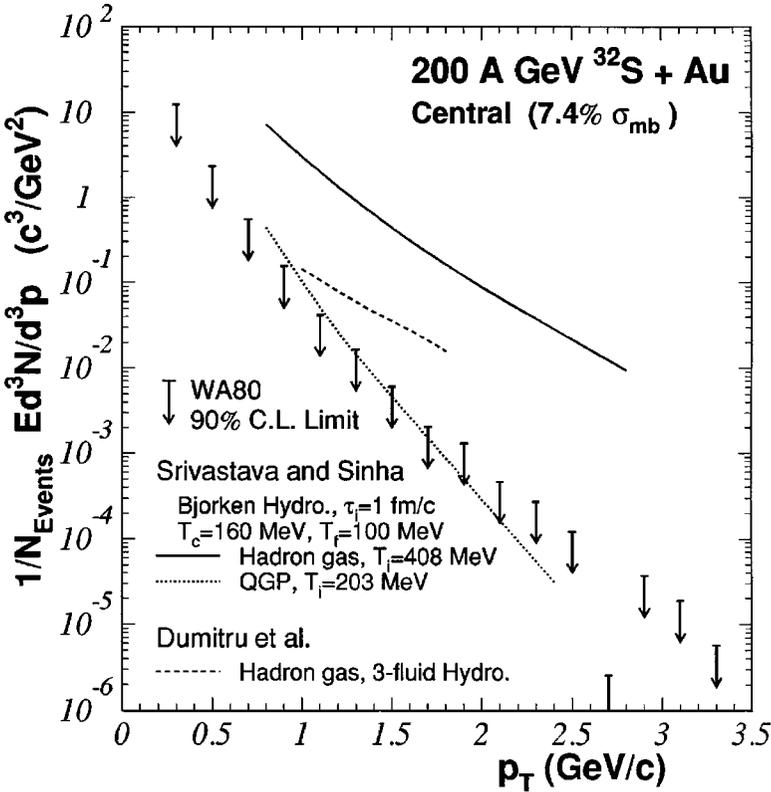


Figure 7 Upper limits on direct photon production in central S+Au collisions at $\sqrt{s_{NN}} = 19$ GeV as measured by the WA80 experiment (63). The bars at the foot of each arrow mark the 90% C.L. upper limit on the differential yield per event (the lengths of the arrows have no significance).

in Section 3.4.1. None of the lighter beam experiments observed a positive signal for direct photons, and so were able to report only upper limits.

3.2.1.1. HELIOS/NA34 The HELIOS/NA34 experiment investigated (59) direct photon production from p , O and S beams on Pt and W targets. Photons were measured through the conversion method using a thin (5.7% of a radiation length) iron plate downstream from the target followed by a charged particle tracking system, covering the rapidity range $1.0 < y < 1.9$ in the lab. The π^0 spectrum was assumed to be equal to the π^- spectrum, which was measured using negative tracks in the same experiment. The ratios of heavier mesons to π^0 s were taken from ISR results, with conservative errors; the dominant contribution to the uncertainty on decay photons was stated to be uncertainty on η production.

Inclusive photon spectra, normalized to the total number of π^0 s, were seen to be consistent with the predicted spectrum of decay photons over a range $0.1 < p_T < 1.5$ GeV/ c for all beam-target combinations. The ratio $r_\gamma \equiv \gamma_{\text{All}}/\pi^0$, defined as integral of all photons with $p_T > 0.1$ GeV/ c (and $p_T > 0.6$ GeV/ c in one case) normalized to the number of pions, was reported for different collision centralities and was always consistent with the result calculated for decay photons. Though they were not reported, it would be possible from knowledge of the errors (4–11% statistical and 9% systematic on the photon data, and 9% on the decay estimate) to calculate upper limits on the integral of the direct photon spectrum from this comparison. However, it should be realized that any estimate of (or constraint on) the integral of a steeply falling p_T spectrum is, in effect, information about only the lowest p_T end of that spectrum.

3.2.1.2. WA80 O beam The WA80 experiment measured photons, π^0 s, and η s simultaneously using a highly segmented lead-glass EM calorimeter. High-energy photons deposit their energy in the form of an EM shower spread in a regular pattern across several calorimeter modules, from which the photon's energy and position can be reconstructed. This reconstruction is complicated in heavy-ion collisions because the high phase-space density of particles leads to significant overlapping of showers, which must then be separated through a clustering algorithm. Stable hadrons will also create hadronic showers in the calorimeter, which must be distinguished from photon-induced EM showers. Discrimination against hadrons is possible based on the shower shapes and with a charged particle veto hodoscope used in WA80; but some remaining hadron contamination, particularly from antineutrons, has to be subtracted away along with the decay photons.

The WA80 experiment (60) reported on a search for direct photons from collisions of p and O beams on C and Au targets, using the in-situ measurements of photons and π^0 s at all p_T , and the η/π^0 ratio in a single p_T bin. The γ/π^0 ratio was measured as a function of p_T in $0.4 < p_T < 2.4$ GeV/ c , for all four beam-target combinations and in separate peripheral and central samples for O+Au. The ratio was systematically higher than γ/π^0 predicted for decay photons in O+C, O+Au and central O+Au, but the excesses were not significant compared to the overall uncertainties. For O+Au collisions, upper limits were set on the direct photon spectrum and not just on a single integral: at 90% C.L. direct photons constituted no more than 15% of the inclusive photon yield at each p_T in the range $0.4 < p_T < 2.4$ GeV/ c .

3.2.1.3. CERES/NA45 The CERES/NA45 experiment carried out an ambitious program of measuring e^+e^- pairs produced in primary $p+A$ and $A+A$ interactions, by identifying and tracking the e^\pm using two ring-imaging Cherenkov (RICH) detectors with an axisymmetric magnetic field in between. Photons which convert inside the primary target are detectable as e^+e^- pairs with vanishing initial opening angle, and a search for direct photon production in S+Au collisions was carried out (61).

The p_T spectrum of photons over the range $0.4 < p_T < 2.0$ GeV/c, and within $2.1 < y < 2.65$, was examined, normalized to the pseudorapidity density of charged particles $dN_{ch}/d\eta$. The π^0 spectrum was inferred in two stages: the shape of the spectrum was constrained to describe the shapes of pion spectra (charged and neutral) from several other S+Au experiments, including the effects of decays of heavier particles; and the magnitude of the π^0 spectrum, relative to $dN_{ch}/d\eta$, was extrapolated from the same ratio in $p+p$ collisions, including a correction for additional baryon stopping in A+A. No evidence was seen for an excess of photons above the calculated decay spectrum. An upper limit was reported on the spectrum integrated over $0.4 < p_T < 2.0$ GeV/c, that at a 90% C.L. direct photons constituted no more than 14% of the integral of inclusive photons in central S+Au collisions. Note that, as mentioned above regarding the HELIOS results, a measurement of such an integral is equivalent to a measurement of the spectrum at the low end of the p_T range.

Additionally, another limit on thermal direct photon production in CERES was derived based on a particular model of thermal production. Since thermal photons are created in the re-scatterings of produced particles, be they quarks and gluons or hadrons, one might expect that the rate of photon production in a volume at some time would go as the square of the density of produced particles in that volume at that time. One might then expect—very naively—that the final-state number of thermal photons would vary as the square of the final-state number of charged particles. If this quadratic relation is assumed to hold perfectly, then the observed relation between the number of photons and N_{ch} as a function of centrality in CERES can be used to set a limit on thermal photons as a fraction of all photons. Such a limit is somewhat lower than the limit from the subtraction method, but it is difficult to know exactly how to use this information in more realistic models.

3.2.1.4. WA80 S beam The WA80 experiment continued its investigations with an expanded calorimeter array to measure photon π^0 and η production in S beam collisions. A preliminary result (62) indicating a significant excess in the γ/π^0 ratio above the decay expectation in central S+Au collisions generated a considerable amount of excitement. However, the final result, measured with the double ratio shown in Figure 5, showed no significant excess of inclusive over decay photons in peripheral or central collisions. Upper limits on the direct photon yield in central S+Au collisions were reported as a function of p_T over $0.5 < p_T < 2.5$ GeV/c; see Figure 7.

In a historical aside, it is interesting to note that the information of upper limits on a direct photon spectrum was not always received in a completely rigorous fashion. At first, some researchers derided limits as indicating only that direct photon results were “missing,” but that view was belied when the WA80 upper limits were shown to provide a significant constraint on theoretical models (see Section 3.3). Others were unused to dealing with having information presented in the form of single-sided limits. For example, in a thorough and well-regarded 1998 review (64) of the field, the state of direct photon measurements in S+Au were described

by the sentence, “Within the reported systematic errors the results of the WA80 and the CERES/NA45 collaborations are compatible with each other.” While true, the statement is tautological: two upper limits (or two lower limits) on the same quantity, or on overlapping quantities, are always compatible with each other.

3.2.2. HEAVY-NUCLEUS BEAMS The CERN-SPS fixed-target heavy ion program continued with ^{208}Pb beams, which were used by upgraded versions of most of the lighter beam experiments. The WA98 experiment, descended from WA80, carried out the same program of measuring photons π^0 's and η 's in Pb+Pb collisions using a highly-segmented lead-glass calorimeter together with a charged particle veto shield. Spectra were examined over $0.5 < p_T < 3.5$ GeV/ c , and the inclusive photons were seen to exceed decay photons at the level of about 2σ in the range $2.0 < p_T < 3.5$ GeV/ c . Thus, for the first time ever in heavy-ion collisions, a direct photon p_T spectrum was reported (58, 65), and not just single-sided limits. These results are shown in Figure 8, and their significance is discussed in Section 3.3.

3.3. Theoretical Interpretation

In the early days of high-energy heavy-ion collisions the first diagnostic goal was simply to show that a multi-collisional, plausibly thermalized state had actually been formed in a collision, and it was hoped that any observation of thermal direct photons would be strong evidence for such a state. So the negative, limits-only results from the early lighter beam experiments were met with some disappointment. But they did spur the development of much more rigorous treatments for diagnosing the states in a collision through direct photons.

To connect a model of a collision to the direct photon spectrum one must, in principle, include all the sources listed in Section 2. A typical calculation would come in two parts: (a) predicting the pQCD prompt photons through standard techniques and (b) predicting the thermal radiation from the QGP and HG phases of matter in the collision as it evolves. Figure 9 shows the result of such a calculation (37), with these three components¹⁵ and their sum. The features shown here are typical: The HG radiation dominates the total number of photons due to that phase's larger volume and longer lifetime; but the QGP radiation dominates at higher photon energy due to its higher temperature, reflecting the $e^{-E/T}$ factor in rate expressions like Equations 3 and 4. The prompt component has even fewer photons, but owing to its flatter power-law shape it will come to dominate at the highest photon energies.

The framework for calculating the thermal components is typically a hydrodynamical description of the collision and its evolution. These models are fast evolving; we reprise only the barest details here. Interested readers can consult (66)

¹⁵A full calculation would also include also pre-equilibrium and parton-medium-interaction photons; but these sources are typically not included other than by authors highlighting specific models for them.

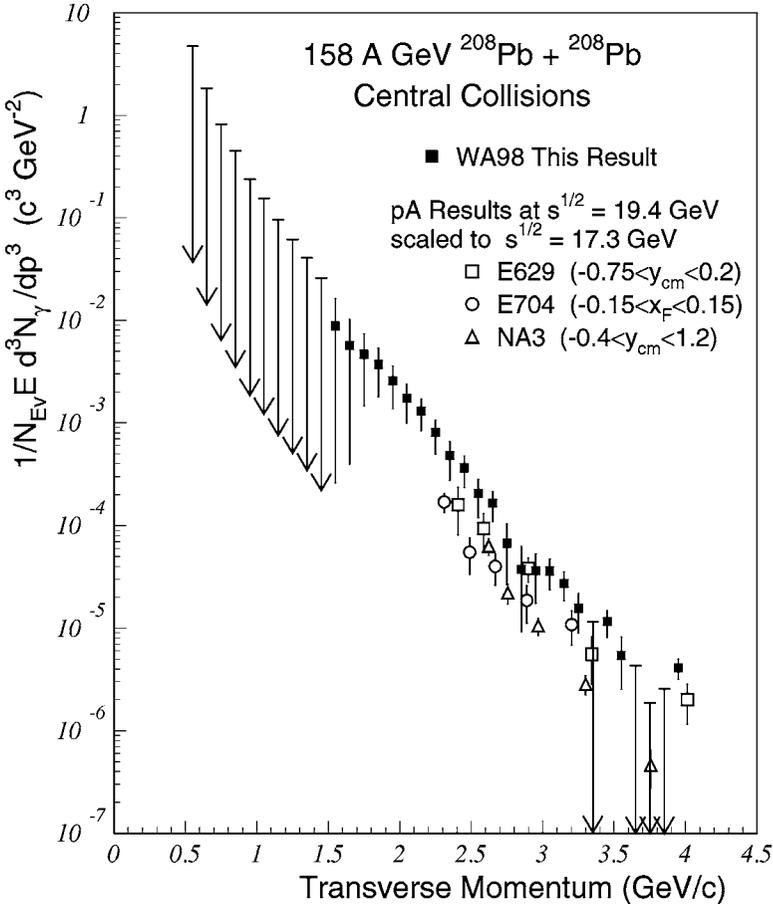


Figure 8 Spectrum of, and upper limits on, direct photon production in central Pb+Pb collisions at $\sqrt{s_{\text{NN}}} = 17$ GeV as measured by the WA98 experiment (65). The points mark the differential yield per event, with standard errors. The bars at the foot of each arrow mark the 90% C.L. upper limit on the yield (the lengths of the arrows have no significance). The hollow points are direct photon yields from elementary collisions at similar energies, which have been re-scaled to the Pb+Pb beam energy and also adjusted via binary scaling to the equivalent yield in central Pb+Pb collisions.

for more detail. The basic ingredients are an initial condition of energy density as a function of position at the time of initial local thermalization, and an equation of state (EOS) relating pressure and temperature to energy density, including any phase transitions (these are typically taken from lattice QCD, for example). The density is then evolved according to energy-momentum conservation, giving each point in space-time a local fluid velocity and a temperature in the local fluid's

rest frame. The radiation from each fluid element can then be computed in its rest frame using a form like Equation 4 and boosted to a common observer's frame to be tallied. The initial conditions are in principle unknown, but can be constrained by comparison to hadron data such as the distribution of final-state transverse energy.

The simple goal of these studies, when compared to data, is to constrain some property of the thermalized state, for example, to show that its density and/or temperature were above the threshold for QGP production. The more ambitious goal is to constrain the equation of state and so learn the properties of hot, thermal QCD matter. A typical exercise is to run a calculation with a few different sets of EOS and initial conditions, and see if any can be ruled out. Less common, but more valuable, are parametric studies exploring which regions of parameter space can, and cannot, be ruled out by comparison to the data.

3.3.1. S BEAM RESULTS A common choice for two EOSs is one with a QGP state as described by lattice QCD and a phase transition at $T = T_C$ to a hadron gas, and one with no phase transition and a pure hadron gas at all temperatures. The hope is, that a no-transition scenario being ruled out can provide evidence for a QCD phase transition taking place in the collision. An example is visible in Figure 7, where photons from these two scenarios calculated by Srivastava & Sinha (67) are compared to the WA80 upper limits on direct photons in S+Au. The pure HG scenario used here incorporated a limited number of hadron species, and this low number of degrees of freedom combined with the assumed initial energy densities resulted in a very high temperature, whose photon rates can easily be ruled out by the data. The scenario with the initial QGP, having a large number of degrees of freedom, is initially cooler and its photons are in much better agreement with the data. A similar conclusion was drawn from the calculations of Dumitru et al. (68).

This comparison does not rule out a pure HG scenario so much as it rules out an HG with artificially few degrees of freedom. Srivastava and Sinha revisited (69) the comparison with S+Au data using an HG incorporating all hadrons with masses below $2.5 \text{ GeV}/c^2$ and concluded that the photons from a no-transition scenario with this richer HG¹⁶ were similar to those from the QGP scenario and were consistent with the data. The implication, noted also by Sarkar et al. (70), was that a blunt transition/no transition distinction could not be drawn from the S+Au data.

A more supple analysis by Sollfrank et al. (72) concluded that WA80 limits ruled out a temperature higher than 250 MeV for the initial phase of a S+Au collision, regardless of its composition. This "cool" initial state, in turn, implies a minimum on the number of degrees of freedom opened in the initial state and so allows us to make contact with the thermodynamics of Figure 3. Inasmuch as a high number of initial degrees of freedom can be taken as evidence for a phase

¹⁶It was noted, however, that these "full" hadron gases have number densities exceeding several/fm³ at temperatures of 200 MeV and above, which may be unphysical; see also Bebie (71).

transition, this evidence gets stronger when fewer photons are observed, in an interesting inversion of the original presumption for direct photons.

3.3.2. PB BEAM RESULTS The positive observation (that is, both upper and lower limits) of direct photons in Pb+Pb reported by WA98 triggered a new, expanded round of theoretical interpretation. No single conclusion has emerged at present, and the sophistication in the treatment of various elements varies widely in different calculations. We will step through some of the main conclusions from a sample of authors.

Dumitru et al. (76) examined prompt pQCD production as the source of photons in Pb+Pb at high $p_T > 3$ GeV/ c , particularly the evidence for k_T broadening. As measured by the parton intrinsic transverse momentum parameter $\langle k_T^2 \rangle$, they found the WA98 data best fit by values in the range 1.8–2.4 GeV²/ c^2 , which is up to 1 GeV²/ c^2 larger than a standard value of 1.3 GeV²/ c^2 called for in $p+p$ photon data. They also concluded that prompt photons could not account for the spectrum at $p_T < 2.5$ GeV/ c , and so a thermal component is required. Steffen and Thoma (39) examined thermal production with a small variety of initial state assumptions, including maximum temperatures up to 235 MeV, and drew the complementary conclusion that a prompt component was necessary to explain the data above $p_T > 2$ GeV/ c .

Srivastava and Sinha (77) varied the initial thermalization time τ_0 and found a good fit to the data with a short $\tau_0 = 0.2$ fm/ c for an initial QGP phase with a correspondingly high temperature of 335 MeV. Huovinen et al. (73) explored different thermalization times and different longitudinal hydrodynamical models, and concluded that the data could be fit with either a scenario with a phase transition or with a hadron-gas-only scenario, but noted that in the latter case the required maximum initial temperature is well above T_C .

Gallmeister et al., motivated by dilepton data, exhibited (74) a fit to the WA98 data with a prompt source plus a simplified, single-temperature, spherically expanding thermal source. Alam et al. (75) incorporated a nonzero initial radial expansion profile and were able to fit the WA98 data with both a QGP initial state and a chirally-symmetric HG initial state, both at modest initial temperatures very close to 200 MeV.

Turbide et al. (37) examined the WA98 photon data as part of their study of hadron gas radiation. Their baseline QGP/HG scenario with a thermalization time of $\tau_0 = 1.0$ fm/ c , initial temperature of 205 MeV, and no additional nuclear k_T broadening underpredicted the data in the range $p_T > 1.5$ GeV/ c . They concluded that the gap could be filled either by shortening the thermalization time and increasing the initial temperature to 270 MeV, or by adding additional broadening of $\langle \Delta k_T^2 \rangle = 0.3$ GeV²/ c^2 to the prompt production.

What emerges is a kind of rough parameter space in which one might vary (a) additional parton nuclear broadening; (b) initial thermalization times, and hence initial temperatures; (c) a QGP/HG versus pure HG EOS; and possibly (d) an initial transverse flow profile. Different authors have shown that various points in this

parameter space can reproduce the WA98 data, but what's missing is a thorough parametric study detailing which segments of parameter space can be ruled out by the data. Such a study has the potential to constrain the thermodynamics of QCD in the SPS Pb+Pb system, and so may hold great promise. Despite the volume of work done so far, there is still more to be learned from the WA98 direct photon data.

3.4. Other Techniques

3.4.1. GAMMA-GAMMA HBT A different approach to fixing the relative rates of direct vs. decay photons takes advantage of the fact that the spatial distribution of the sources of the two populations differ widely. Direct photons from a hadron collision are all generated from a volume whose size is set by the scale of strong interactions and the size of the nucleus, on the (rough) order of femtometers (fm). The decay photons are born in a volume with typical dimension on the order of the decay length for an EM decay—25 nanometers (nm) for π^0 —which is many orders of magnitude larger. The distributions of single inclusive photons are not sensitive to these shapes, but the distributions of pairs of photons can be. Specifically, the Hanbury-Brown-Twiss (HBT) technique of identical boson interferometry can relate the size and shape of a source distribution to the intensity and shape of a photon pairs distribution as a function of their relative momentum.

The HBT technique was first employed (78) to measure the sizes of distant stars. But over the last several decades it has become very highly developed (79, 80) for reconstructing the shapes of secondary source distributions from high-energy particle collisions, which are smaller than stars by more than 20 orders of magnitude. We will give only the briefest outline of the technique here, highlighting the relevance to measuring direct photons (see References 81–84 for more details on this application).

Figure 10 illustrates the basic idea. A source is incoherently radiating some kind of particle from some finite region of space. We consider the case that two identical particles with very similar energies are emitted very close together in time from two separate points on the source and are then observed in coincidence in two detectors D_1 and D_2 . Classically we would expect that the rate of pair coincidences would follow from the rates of singles in each detector, $\Gamma(D_1, D_2) \propto \Gamma(D_1)\Gamma(D_2)$.

Quantum mechanics alters this picture, however. For a coincidence to occur, there are two possible combinations of paths that the particles can take from the source points to the detectors, shown as solid and dashed lines in Figure 10. Because the two particles of the pair are identical, these two possibilities are indistinguishable—they have identical initial conditions and identical final states—and so we must add their quantum mechanical amplitudes in order to predict the rate of coincidence events. We do not recap the mathematics here, but the basic result is that for identical boson pairs the coincidence rate will be enhanced for pairs with small relative momentum Δp , specifically when Δp is on the order of (or smaller

than \hbar/R where R is the size of the source. This is depicted schematically in the middle row of Figure 10. First we scale the coincidence rate by the product of the singles rates to form a correlation function $\Gamma(D_1, D_2)/\Gamma(D_1)\Gamma(D_2)$ as a function of pair Δp and normalize it to 1 at high $\Delta p \gg \hbar/R$. The correlation function will then show a peak for $\Delta p < \sim \hbar/R$ and approach a value of 2 as $\Delta p \rightarrow 0$ for pairs of spinless bosons.¹⁷ For photons, the correlation function peaks at a value of $3/2$ —the strength of the enhancement is reduced by half because two photons act as identical particles only if they have the same polarization, and this will be true of only half the pairs.

The spatial source distribution of photons emerging from a high-energy hadron collision, however, has pieces at two length scales: a core on the fm scale, from which the direct photons are produced, and a halo on the nm scale, from which the decay photons are produced. The photon-photon correlation function for this source will also have features on two scales. Photon pairs where both are from the core, i.e., direct-direct pairs, will show an enhanced coincidence rate for momentum differences on the order of $\Delta p \sim \hbar/\text{fm} \sim 100 \text{ MeV}/c$; whereas the other pairs, the direct-decay and decay-decay, will show enhancement on the much smaller scale of $\Delta p \sim \hbar/\text{nm} \sim 100 \text{ eV}/c$. The resulting correlation function is depicted schematically in the bottom row of Figure 10, rising to its ideal height of $3/2$ at $\Delta p = 0$.

In practice the narrower feature is never seen in high-energy collisions, for two reasons. First, pairs of hard gamma rays with these tiny values of Δp would have an opening angle so small that it would be extremely difficult to detect them separately. Second, the phase space for pairs with tiny Δp is so small that they would essentially never be observed in a typical experiment. So in practice we would expect the measured photon-photon correlation function to have a single-peak shape over Δp , but rising to an apparent peak value (well) below $3/2$.

If f is the fraction all photons which are direct, then the direct-direct pairs are a fraction f^2 of all pairs. The low- Δp enhancement will be reduced by exactly this fraction, and the wide feature in the correlation function will (in the ideal case) have a peak value of $1 + f^2/2$ as indicated in Figure 10. So the peak value of the correlation function, combined with a measure of the inclusive photon rate, can in principle yield a measurement of the direct photon rate at that energy.

Extracting direct photon yields via photon-photon HBT has great appeal because it does not require any measurement of π^0 's or other hadrons; but it is experimentally very challenging. First results for direct photon yields extracted this way at low p_T in Pb+Pb were recently reported by the WA98 experiment [86, see also Reference 85 (Peressounko)]. The data are shown in Figure 11 and lie

¹⁷This is true in an idealized situation, while the actual peak values are affected by whether the source is truly incoherent and by residual strong and EM interactions between the particles. Because photons do not suffer the latter, the case of photon-photon HBT will be closer to ideal than hadron-hadron HBT.

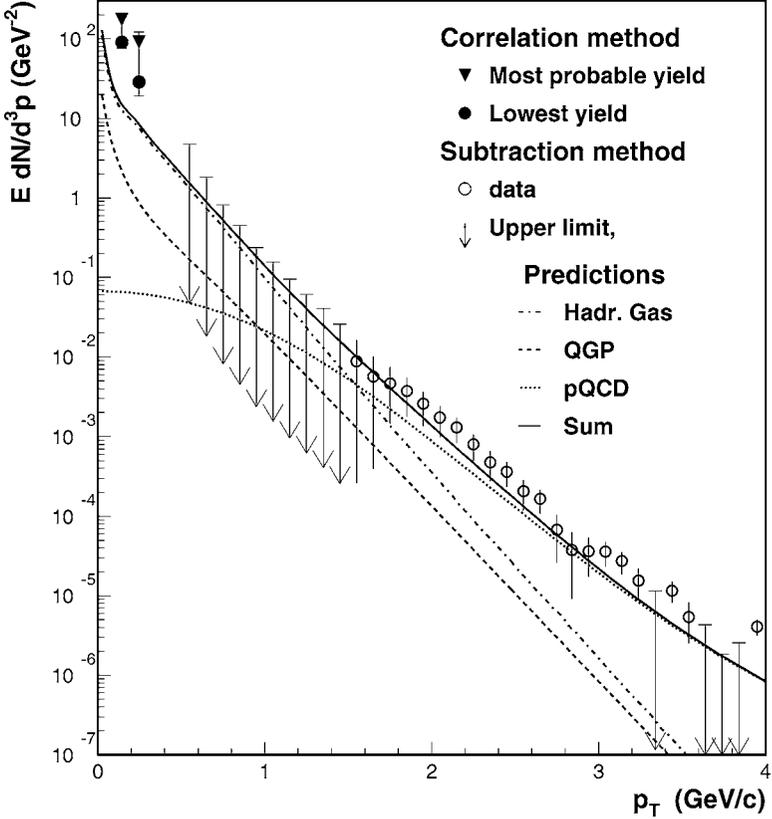


Figure 11 Direct photon spectrum from Pb+Pb collisions at CERN-SPS ($\sqrt{s_{NN}} = 17$ GeV), now extended to lower p_T with the photon-photon HBT technique (86). The curves are direct photon rates as calculated (37) for different sources. The solid line is the sum, which is in reasonable agreement with the higher- p_T results but badly underpredicts the rate at lower p_T .

well above a typical theoretical expectation. We would expect that direct photons at these low p_T would be primarily hadron-gas radiation, and the shapes of the photon-photon correlation functions indicate a source size that would be consistent with the later, expanded stages of a heavy-ion collision. The implications of the low- p_T yield and the correlation shapes for HG dynamics are presently under very active investigation (9, 85, 88).

3.4.2. PHOTON-BY-PHOTON IDENTIFICATION The techniques described above, statistical subtraction and photon-photon HBT, can give information only on the average yield of direct photons over a large event sample. In order to make a hadron-direct photon pair measurement, for example to measure photon-tagged jets

(Section 2.1) or to search for medium-induced photon jet fragments (Section 2.4), it is necessary to identify direct photons individually, on a photon-by-photon basis, with reasonable efficiency.

3.4.2.1. Isolation As mentioned in Section 1, a standard approach in elementary particle collisions is to measure photons that have no accompanying energy, either hadronic or EM, within some angular range. Such an object, defined as an isolated photon, is reasonably assured to be a direct photon, at least at very high momentum, because the parent hadron of a decay photon would be part of a jet and other hadrons from the jet would be nearby at close angles. Isolated photons can be measured in calorimeters with relatively coarse segmentation and lower energy resolution, since fine precision on their position and energy is not required. Treating isolated photons theoretically is complicated, because the definition of the isolation cut and the properties of jet fragmentation enter into calculations of their rates, and the definition will exclude fragmentation/bremsstrahlung photons.

Although isolated photons are quite standard as observable objects in elementary collisions, it is hopeless to search for them from high-energy nuclear collisions because the high densities of produced particles will essentially always deposit some energy within any standard isolation cone around a photon.

3.4.2.2. Sister-tagging An alternative approach to individual direct photon identification that might be tractable in heavy-ion collisions is to veto a candidate photon if an accompanying “sister photon” is found, which indicates that both could have been produced in an hadronic decay. Operationally one would, for each candidate photon, run through all pairs with other photons in the event and check to see if the invariant mass $m_{\gamma\gamma}$ of any pair was consistent with the mass of a hypothetical parent π^0 or η . Though the term is not in common use, this method could reasonably be termed sister-matching or sister-tagging. The approach requires an EM calorimeter with sufficient energy and position resolution to provide good accuracy in $m_{\gamma\gamma}$ in order to minimize efficiency losses from accidental rejections, though these will always be present at some level. Also, some false positive identification of decay photons as direct will always occur at some level, for two reasons: Not all potential parent decays can be identified in a two-photon channel; and the opening angle between two photon daughters can, regardless of parent energy, always¹⁸ be as large as 180° , i.e., perfectly back-to-back, so without true 4π coverage the sister always has some chance to be lost.

Sister-tagging is also used to measure inclusive direct photons in elementary high-energy collisions (16, 89), and direct photons identified this way are easier to treat theoretically than are isolated photons. Its utility for identification of individual direct photons in heavy-ion collisions has not yet been exhibited, but work continues actively.

¹⁸This interesting property is unique to two-body decays with massless daughters.

3.5. Collider Experiments: First Results and Future Prospects

3.5.1. FIRST RESULTS AT RHIC Heavy-ion collisions at RHIC have been examined by four experiments starting in the year 2000, and the wealth of data accumulated so far has provided strong evidence (90) that a dense, thermalized state is formed early in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Many questions are not yet answered, such as whether deconfinement can be shown to have occurred and what might be the mechanism through which fast initial thermalization is achieved. The measurement of direct photons can offer unique information in furthering this investigation.

RHIC A+A collisions are expected to provide a rich field for direct photon physics, because all of the components we have discussed—prompt, pre-equilibrium, thermal, and parton-medium-interaction induced—have been predicted to appear at significant, observable levels. There are also several advantages to measuring and interpreting direct photons at RHIC compared with the SPS. On the theoretical side, the spectrum of π^0 s and prompt direct photons measured in $p+p$ collisions at RHIC by the PHENIX experiment (91, 92) have been well-reproduced by next-to-leading order (NLO) pQCD calculations. The agreement is good even down to the low p_T ranges, where disagreement was seen between data and theory (12, 16) for lower beam energies, suggesting that the prompt component is under reasonable control at RHIC. On the experimental side, the depletion of high- p_T π^0 s and other hadrons in A+A collisions greatly reduces the yield of decay photons at RHIC, which greatly eases the task of extracting direct photons (see Figure 6).

The PHENIX experiment has recently reported (93, 94) spectra of direct photons in Au+Au collisions, for a variety of centralities, obtained with the statistical subtraction method. The direct photon spectra are shown in Figure 12, and compared to the NLO calculation for $p+p$ collisions at the same $\sqrt{s_{NN}}$, adjusted according to binary scaling for that centrality. Interestingly, the binary-scaled NLO calculation is consistent with the direct photon data at all centralities, suggesting, at least at first glance, that there are no large enhancements from other sources compared to the prompt source. The fact that direct photons obey point-like scaling whereas high- p_T hadrons do not is reflected in the great increase in the γ/π^0 ratio with centrality seen in Figure 6. This provides a dramatic confirmation that the depletion of high- p_T hadrons, i.e., jet quenching, is due to a final-state effect on scattered partons, as discussed in Section 2.1.

What the implications might be of small yields from the other sources, exactly how much room remains between the PHENIX data and the calculated prompt yield, and how reliable a prompt calculation can be are all questions that we expect to be investigated with great interest in the near future. In a very recent paper (95) d'Enterria and Peressounko employed a hydrodynamical model of RHIC Au+Au collisions, as constrained by RHIC hadron data, to calculate a thermal direct photon

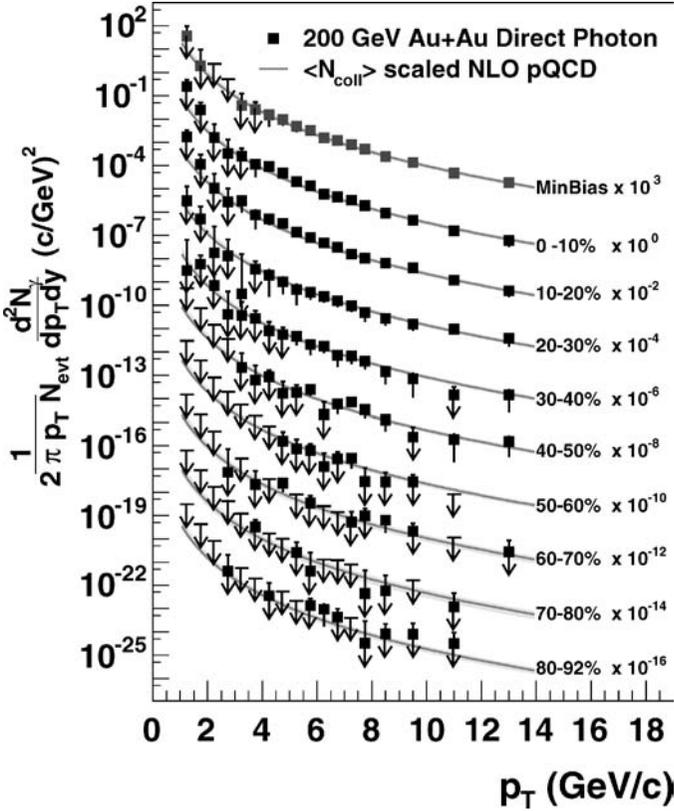


Figure 12 Spectrum of, and upper limits on, direct photon production in Au+Au collisions at RHIC, $\sqrt{s_{NN}} = 200$ GeV, as measured by the PHENIX experiment (93). The points mark the differential yield per event, with standard errors. The bars at the foot of each arrow mark the 90% C.L. upper limit on the yield (the lengths of the arrows have no significance).

spectrum. They reported that the thermal photons should dominate the prompt for $p_T < 3$ GeV/c, and that the sum spectrum is consistent with the PHENIX data for central and peripheral collisions over the whole p_T range down to 1 GeV/c. They also exhibit a relation between the local slope of the thermal photon spectrum and the initial temperature of the system, calculated through the hydrodynamical model, and propose a method by which temperature information can be combined with entropy density information (from the measured charged particle multiplicity) to constrain the number of degrees of freedom in the initial state. This approach is similar in spirit to that of Sollfrank et al. (72), who combined temperature and energy density to investigate degrees of freedom.

Many experimental advances are expected soon in direct photons at RHIC. The direct photon spectra discussed here were derived from only a small fraction of the now-available A+A and $p+p$ data, so great improvements in statistical accuracy should be possible. Also, the very large $\gamma^{Direct}/\gamma^{Decay}$ ratio, as a result of jet quenching of mesons, means that identifying individual direct photons is much easier, which greatly facilitates measurements like γ^{Direct} -tagged jets and γ^{Direct} -hadron correlations. All in all, there is every reason to expect that direct photons will soon be a unique and important part of the investigation of RHIC collisions.

3.5.2. FUTURE PROSPECTS AT LHC Beyond RHIC, the next stage on the beam-energy frontier in heavy-ion collisions is the CERN-LHC. The LHC will start providing Pb+Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV before the end of this decade, an increase by a larger factor over RHIC ($\times 27$) than RHIC was over SPS ($\times 10$). Unlike RHIC, the LHC will be primarily dedicated to high-energy particle physics. As currently planned, the LHC heavy-ion program will run for one month out of the year, with the balance going for $p+p$ collisions at top energy of $\sqrt{s} = 18$ TeV.

The role of direct photons in nuclear collisions at the LHC is discussed in great detail in the “Photon Physics” chapter of the *CERN Yellow Report on Hard Probes in Heavy Ion Collisions at the LHC* (96). We touch here on some general considerations to keep in mind.

Higher energy scales in LHC collisions will change the balance between different sources of direct photons, but it is not yet clear how. The prompt component will have a much greater yield and a flatter spectrum from LHC collisions than at RHIC. But initial temperatures are also expected to be higher—numbers between 300 MeV and 800 MeV have been mentioned—which could still lead to a QGP thermal component being dominant in some p_T range. Various authors have reached different conclusions depending on the particulars of their assumptions.

Prompt photons with p_T s of few–10 GeV/ c will probe much lower Bjorken x 's in the incoming nuclei, and so deepen the investigation of possible coherent/saturated gluon initial states such as the CGC. At the same time, high energy jets with many 10s of GeV will be produced at significant rates at LHC, and so photon-jet, or possibly Z-jet, tagging will be of great importance in studying possible modifications to their fragmentation.

Finally, pre-equilibrium processes are likely to be of greater significance at LHC. Parton scatterings at LHC are overwhelmingly dominated by gluons—the machine is sometimes described as a “gluon collider”—and so the possibility that early thermalized states may be out of chemical equilibrium, gluon rich but quark poor, must be kept in mind. In our opinion, the as-yet-unknown physics of the pre-equilibrium stage may be the next great subject of study in heavy-ion collisions, yielding greater surprises than the study of the thermal QGP. And whatever happens in those very early stages, the same general principle applies, that direct photons may carry and preserve the best diagnostic information about them.

The Annual Review of Nuclear and Particle Science is online at
<http://nucl.annualreviews.org>

LITERATURE CITED

1. Feinberg EL. *Nuovo Cim.* A34:391 (1976)
2. Shuryak EV. *Phys. Lett.* B78:150 (1978)
3. McLerran LD, Toimela T. *Phys. Rev. D* 31: 545 (1985)
4. Peitzmann T, Thoma MH. *Phys. Rep.* 364: 175 (2002)
5. Gale C, Haglin KL. arXiv:hep-ph/0306098
6. Alam J, Sinha B, Raha S. *Phys. Rep.* 273: 243 (1996)
7. Aurenche P. arXiv:hep-ph/0201011
8. Gale C. *Nucl. Phys.* A698:143 (2002)
9. Rapp R. *Mod. Phys. Lett.* A19:1717 (2004)
10. d'Enterria D. *Phys. Lett.* B596:32 (2004)
11. McLerran LD. *Lect. Notes Phys.* 583:291 (2002)
12. Aurenche P, et al. *Eur. Phys. J.* C9:107 (1999); Aurenche P, et al. *Phys. Rev. D* 55: 1124 (1997)
13. Wang XN, Huang Z. *Phys. Rev. C* 55:3047 (1997)
14. Srivastava DK, Gale C, Awe TC. *Phys. Rev. C* 67:054904 (2003)
15. Arleo F, Aurenche P, Belghobsi Z, Guillet JP. *J. High Energy Phys.* 0411:009 (2004)
16. Apanasevich L, et al. (Fermilab E706 Col- lab.) *Phys. Rev. D* 70:092009 (2004)
17. Creutz M. *Phys. Rev. D* 15:1128 (1977)
18. Karsch F. *Lect. Notes Phys.* 583:209 (2002)
19. Laermann E, Philipsen O. *Annu. Rev. Nucl. Part. Sci.* 53:163 (2003)
20. Petreczky P. *Nucl. Phys. Proc. Suppl.* 140: 78 (2005)
21. Matsubara T. *Prog. Theor. Phys.* 14:351 (1955)
22. Rapp R, Wambach J. *Adv. Nucl. Phys.* 25:1 (2000); Rapp R. *J. Phys.* G31:S217 (2005)
23. Weldon HA. *Phys. Rev. D* 28:2007 (1983)
24. Gale C, Kapusta JI. *Nucl. Phys.* B357:65 (1991)
25. Thoma MH. *Phys. Rev. D* 51:862 (1995)
26. Kajantie K, Ruuskanen PV. *Phys. Lett.* B121:352 (1983)
27. Kapusta JI, Lichard P, Seibert D. *Phys. Rev. D* 44:2774 (1991). Erratum. *Phys. Rev. D* 47:4171 (1993)
28. Baier R, Nakkagawa H, Niegawa A, Redlich K. *Z. Phys. C* 53:433 (1992)
29. Braaten E, Pisarski RD. *Nucl. Phys.* B337:569 (1990); Braaten E, Pisarski RD. *Phys. Rev. D* 42:2156 (1990); Braaten E, Pisarski RD. *Phys. Rev. D* 45:1827 (1992)
30. Aurenche P, Gelis F, Kobes R, Zaraket H. *Phys. Rev. D* 58:085003 (1998); Aurenche P, Gelis F, Kobes R, Zaraket H. *Phys. Rev. D* 60:076002 (1999)
31. Arnold P, Moore GD, Yaffe LG. *J. High Energy Phys.* 0111:057 (2001)
32. Arnold P, Moore GD, Yaffe LG. *J. High Energy Phys.* 0112:009 (2001)
33. Karsch F, et al. *Phys. Lett. B* 530:147 (2002)
34. Brown GE, Lee CH, Rho M, Shuryak E. *Nucl. Phys.* A740:171 (2004)
35. Karsch F, et al. *Nucl. Phys.* A715:701 (2003)
36. Xiong L, Shuryak EV, Brown GE. *Phys. Rev. D* 46:3798 (1992)
37. Turbide S, Rapp R, Gale C. *Phys. Rev. C* 69: 014903 (2004)
38. Turbide S, Rapp R, Gale C. *Int. J. Mod. Phys.* A19:5351 (2004)
39. Steffen FD, Thoma MH. *Phys. Lett.* B510: 98 (2001)
40. Shuryak EV, Xiong L. *Phys. Rev. Lett.* 70: 2241 (1993)
41. Traxler CT, Vija H, Thoma MH. *Phys. Lett.* B346:329 (1995); Traxler CT, Thoma MH. *Phys. Rev. C* 53:1348 (1996)
42. Gelis F, Niemi H, Ruuskanen PV, Rasanen SS. *J. Phys.* G30:S1031 (2004)
43. Srivastava DK, Geiger K. *Phys. Rev. C* 58: 1734 (1998)
44. Dumitru A, et al. *Phys. Rev. C* 57:3271 (1998)

45. Boyanovsky D, de Vega HJ. *Phys. Rev. D* 68:065018 (2003); Boyanovsky D, de Vega HJ. *Nucl. Phys.* A747:564 (2005)
46. Serreau J. *J. High Energy Phys.* 0405:078 (2004)
47. Jeon SY, Jalilian-Marian J, Sarcevic I. *Phys. Lett.* B562:45 (2003)
48. Fries RJ, Müller B, Srivastava DK. *Phys. Rev. Lett.* 90:132301 (2003)
49. Zakharov BG. *JETP Lett.* 80:1 (2004) [*Pisma Zh. Eksp. Teor. Fiz.* 80:3 (2004)]
50. Srivastava DK, Gale C, Fries RJ. *Phys. Rev. C* 67:034903 (2003); Gale C, Awes TC, Fries RJ, Srivastava DK. *J. Phys.* G30: S1013 (2004)
51. Vogt R. *Phys. Rev. C* 64:044901 (2001)
52. Kartvelishvili V, Kvataadze R, Shanidze R. *Phys. Lett.* B356:589 (1995)
53. Kapusta JI, Wong SMH. *Phys. Rev. D* 62: 037301 (2000)
54. Kapusta JI, Wong SMH. *Phys. Rev. D* 64: 045008 (2001)
55. Aurenche P, et al. *Phys. Rev. D* 65:038501 (2002)
56. Kapusta JI, Wong SMH. *Phys. Rev. D* 65: 038502 (2002)
57. Majumder A, Gale C. *Phys. Rev. C* 65: 055203 (2002)
58. Aggarwal MM, et al. (WA98 Collab.) arXiv:nucl-ex/0006007 *Phys. Rev. C*. Submitted
59. Akesson T, et al. *Z. Phys. C* 46:369 (1990)
60. Albrecht R, et al. (WA80 Collab.) *Z. Phys. C* 51:1 (1991)
61. Baur R, et al. (CERES Collab.) *Z. Phys. C* 71:571 (1996); Kampert KH, et al. *Z. Phys. C* 74:587 (1997); Baur R, et al. *Z. Phys. C* 74:593 (1997)
62. Santo R, et al. (WA80 Collab.) *Nucl. Phys.* A566:61C (1994)
63. Albrecht R, et al. (WA80 Collab.) *Phys. Rev. Lett.* 76:3506 (1996)
64. Bass SA, Gyulassy M, Stocker H, Greiner W. *J. Phys.* G25:R1 (1999)
65. Aggarwal MM, et al. (WA98 Collab.) *Phys. Rev. Lett.* 85:3595 (2000)
66. Huovinen P, Ruuskanen V. *Annu. Rev. Nucl. Part. Sci.* 56:In press (2006)
67. Srivastava DK, Sinha B. *Phys. Rev. Lett.* 73:2421 (1994)
68. Dumitru A, et al. *Phys. Rev. C* 51:2166 (1995)
69. Srivastava DK, Sinha BC. *Eur. Phys. J.* C12:109 (2000). Erratum. *Eur. Phys. J.* C20:397 (2001)
70. Sarkar S, Roy P, Alam J, Sinha B. *Phys. Rev. C* 60:054907 (1999)
71. Bebie H, Gerber P, Goity JL, Leutwyler H. *Nucl. Phys.* B378:95 (1992)
72. Sollfrank J, et al. *Phys. Rev. C* 55:392 (1997)
73. Huovinen P, Ruuskanen PV, Rasanen SS. *Phys. Lett.* B535:109 (2002)
74. Gallmeister K, Kampfer B, Pavlenko OP. *Phys. Rev. C* 62:057901 (2000)
75. Alam J, Sarkar S, Hatsuda T, Nayak TK, Sinha B. *Phys. Rev. C* 63:021901 (2001)
76. Dumitru A, et al. *Phys. Rev. C* 64:054909 (2001)
77. Srivastava DK, Sinha B. *Phys. Rev. C* 64: 034902 (2001)
78. Hanbury Brown R, Twiss RQ. *Philos. Mag.* 45:663 (1954)
79. Lisa M, Pratt S, Soltz R, Wiedemann U. *Annu. Rev. Nucl. Part. Sci.* 55:In press (2005)
80. Baym G. *Acta Phys. Polon.* B29:1839 (1998). arXiv:nucl-th/9804026
81. Srivastava DK, Kapusta JI. *Phys. Rev. C* 48:1335 (1993); Srivastava DK, Kapusta JI. *Phys. Rev. C* 50:505 (1994)
82. Marques FM, et al. *Phys. Rept.* 284:91 (1997)
83. Peressounko D. *Phys. Rev. C* 67:014905 (2003)
84. Bass SA, Müller B, Srivastava DK. *Phys. Rev. Lett.* 93:162301 (2004)
85. Alam J, et al. *Phys. Rev. C* 67:054902 (2003)
86. Aggarwal MM, et al. (WA98 Collab.) *Phys. Rev. Lett.* 93:022301 (2004)
87. Peressounko D, (WA98 Collab.) *J. Phys.* G30:S1065 (2004)
88. Srivastava DK. *Phys. Rev. C* 71:034905 (2005)

89. Bonesini M, et al. (WA70 Collab.) *Z. Phys. C* 38:371 (1988)
90. Arsene I. *Nucl. Phys. A* 757:1 (2005). Adcox K. *Nucl. Phys. A* 757:184 (2005). Back BB. *Nucl. Phys. A* 757:28 (2005). Adams J. *Nucl. Phys. A* 757:102 (2005)
91. Adler SS, et al. (PHENIX Collab.) *Phys. Rev. Lett.* 94:232301 (2005) arXiv:nucl-ex/0503003
92. Adler SS, et al. (PHENIX Collab.) *Phys. Rev. Lett.* 91:241803 (2003)
93. Adler SS, et al. (PHENIX Collab.) arXiv:nucl-ex/0503003
94. Frantz J, (PHENIX Collab.) *J. Phys.* G30:S1003 (2004)
95. d'Enterria D, Peressounko D. arXiv:nucl-th/0503054
96. Arleo F, et al. arXiv:hep-ph/0311131

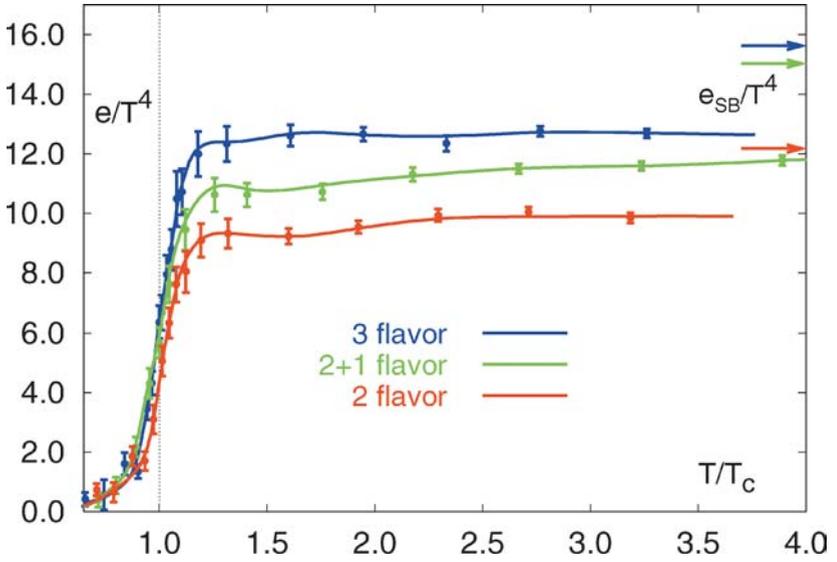


Figure 3 Lattice QCD results for normalized energy density ϵ/T^4 versus temperature, for the assumption of three colors, zero net baryon number, and different numbers of quark flavors (from Karsch (18)). The steep rise suggests a phase transition. The arrows (*right*) mark the Stefan-Boltzmann value for each assumption, i.e., the energy density ϵ_{SB} that an ideal relativistic gas with the corresponding number of particle types would have.

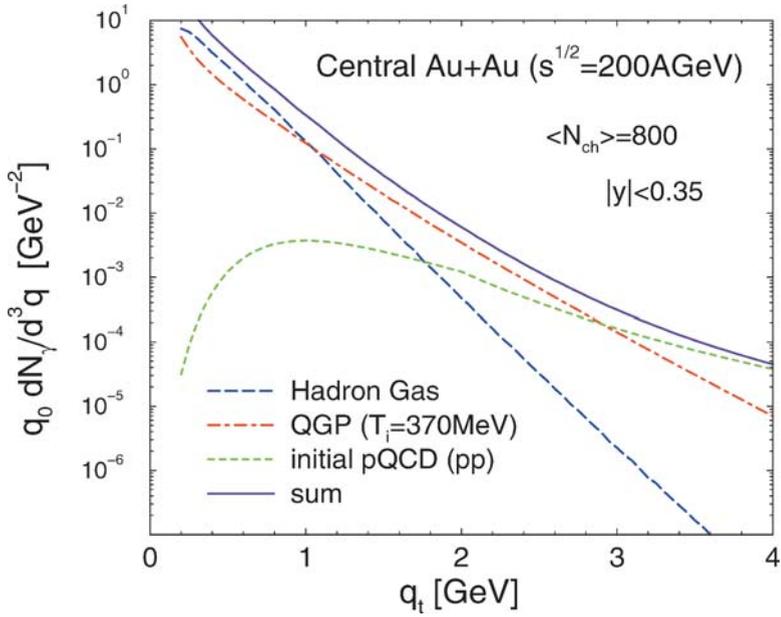


Figure 9 Predicted direct photon spectrum from central RHIC Au+Au collisions, showing the contributions from initial prompt pQCD scatterings and thermal radiation from QGP and hadron gas phases [from Reference (37)].

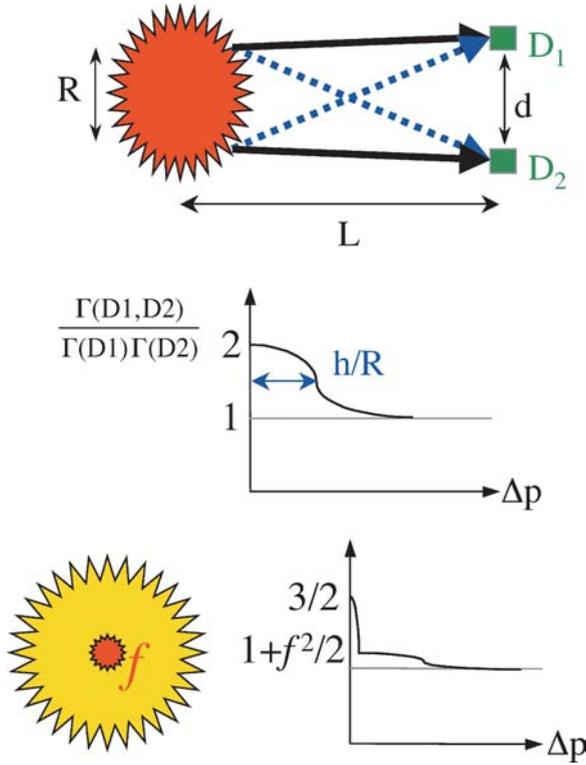


Figure 10 (Upper row) Two possible combinations of paths (solid vs. dotted lines) for two photons from points on a radiating source to detectors D_1 and D_2 . (Middle row) An enhancement in the pair coincidence rate is observed when $R \Delta p < \sim \hbar$, where R is the size of the source and Δp is the relative momentum between the two photons. (Bottom row) When the source has pieces at two different lengths scales, the enhancement pattern exhibits features at two different momentum scales.

CONTENTS

FRONTISPIECE, <i>D.H. Perkins</i>	xii
FROM PIONS TO PROTON DECAY: TALES OF THE UNEXPECTED, <i>D.H. Perkins</i>	1
FUNDAMENTAL NEUTRON PHYSICS, <i>Jeffrey S. Nico</i> <i>and W. Michael Snow</i>	27
TOWARD REALISTIC INTERSECTING D-BRANE MODELS, <i>Ralph Blumenhagen, Mirjam Cvetič, Paul Langacker,</i> <i>and Gary Shiu</i>	71
BLIND ANALYSIS IN NUCLEAR AND PARTICLE PHYSICS, <i>Joshua R. Klein and Aaron Roodman</i>	141
STUDY OF THE FUNDAMENTAL STRUCTURE OF MATTER WITH AN ELECTRON-ION COLLIDER, <i>Abhay Deshpande, Richard Milner,</i> <i>Raju Venugopalan, and Werner Vogelsang</i>	165
LITTLE HIGGS THEORIES, <i>Martin Schmaltz and David Tucker-Smith</i>	229
PHYSICS OF ULTRA-PERIPHERAL NUCLEAR COLLISIONS, <i>Carlos A. Bertulani, Spencer R. Klein, and Joakim Nystrand</i>	271
LEPTOGENESIS AS THE ORIGIN OF MATTER, <i>W. Buchmüller,</i> <i>R.D. Peccei, and T. Yanagida</i>	311
FEMTOSCOPY IN RELATIVISTIC HEAVY ION COLLISIONS: TWO DECADES OF PROGRESS, <i>Michael Annan Lisa, Scott Pratt,</i> <i>Ron Soltz, and Urs Wiedemann</i>	357
SMALL-X PHYSICS: FROM HERA TO LHC AND BEYOND, <i>Leonid Frankfurt, Mark Strikman, and Christian Weiss</i>	403
ASCERTAINING THE CORE COLLAPSE SUPERNOVA MECHANISM: THE STATE OF THE ART AND THE ROAD AHEAD, <i>Anthony Mezzacappa</i>	467
DIRECT PHOTON PRODUCTION IN RELATIVISTIC HEAVY-ION COLLISIONS, <i>Paul Stankus</i>	517
TOOLS FOR THE SIMULATION OF HARD HADRONIC COLLISIONS, <i>Michelangelo L. Mangano and Timothy J. Stelzer</i>	555